SPSP 2020

International Society for Philosophy of Science in Practice

Abstracts of Research Accepted for Presentation* at SPSP 2020

**Due to the global Covid-19 pandemic, the 8th biennial conference of the Society for the Philosophy of Science in Practice, scheduled to take place at Michigan State University in July 2020, was cancelled. Prior to cancellation, conference submissions had been reviewed and the list of accepted research was finalized. This booklet contains abstracts of all research accepted for presentation at SPSP 2020, except in cases where an author preferred that their abstract not be published.*

Table of Contents

I. About SPSP 3
II. Organising Committees 4
III. SPSP Statement on the Covid-19 Pandemic.. 5
IV. Abstracts of Symposia 6
V. Abstracts of Contributed Talks 49

I. About SPSP

Philosophy of science has traditionally focused on the relation between scientific theories and the world, at the risk of disregarding scientific practice. In social studies of science and technology, the predominant tendency has been to pay attention to scientific practice and its relation to theories, sometimes willfully disregarding the world except as a product of social construction. Both approaches have their merits, but they each offer only a limited view, neglecting some essential aspects of science. We advocate a philosophy of scientific practice, based on an analytic framework that takes into consideration theory, practice and the world simultaneously.

The direction of philosophy of science we advocate is not entirely new: naturalistic philosophy of science, in concert with philosophical history of science, has often emphasized the need to study scientific practices; doctrines such as Hacking's "experimental realism" have viewed active intervention as the surest path to the knowledge of the world; pragmatists, operationalists and late-Wittgensteinians have attempted to ground truth and meaning in practices. Nonetheless, the concern with practice has always been somewhat outside the mainstream of English-language philosophy of science. We aim to change this situation, through a conscious and organized programme of detailed and systematic study of scientific practice that does not dispense with concerns about truth and rationality.

Practice consists of organized or regulated activities aimed at the achievement of certain goals. Therefore, the epistemology of practice must elucidate what kinds of activities are required in generating knowledge. Traditional debates in epistemology (concerning truth, fact, belief, certainty, observation, explanation, justification, evidence, etc.) may be re-framed with benefit in terms of activities. In a similar vein, practice-based treatments will also shed further light on questions about models, measurement, experimentation, etc., which have arisen with prominence in recent decades from considerations of actual scientific work.

There are some salient aspects of our general approach that are worth highlighting here.

1. We are concerned with not only the acquisition and validation of knowledge, but its use. Our concern is not only about how pre-existing knowledge gets applied to practical ends, but also about how knowledge itself is fundamentally shaped by its intended use. We aim to build meaningful bridges between the philosophy of science and the newer fields of philosophy of technology and philosophy of medicine; we also hope to provide fresh perspectives for the latter fields.
2. We emphasize how human artifacts, such as conceptual models and laboratory instruments, mediate between theories and the world. We seek to elucidate the role that these artifacts play in the shaping of scientific practice.
3. Our view of scientific practice must not be distorted by lopsided attention to certain areas of science. The traditional focus on fundamental physics, as well as the more recent focus on certain areas of biology, will be supplemented by attention to other fields such as economics and other social/human sciences, the engineering sciences, and the medical sciences, as well as relatively neglected areas within biology, physics, and other physical sciences.
4. In our methodology, it is crucial to have a productive interaction between philosophical reasoning and a study of actual scientific practices, past and present. This provides a strong rationale for history-and-philosophy of science as an integrated discipline, and also for inviting the participation of practicing scientists, engineers and policymakers.

II. Organising Committees

While we did not hold a meeting in 2020, the SPSP Steering and Organizing Committees invested considerable work in preparing for SPSP 2020, including setting up the call for papers, reviewing conference abstracts and finalising the conference programme. Pre-pandemic, the Local Organising Committee had already committed substantial time and resources, including completing extensive fundraising activities, reserving accommodations, booking venues, and coordinating other necessary preparations for the conference and its local hosting. These committees also worked together in making the difficult decisions regarding the unfeasibility of convening the 8th biennial conference during the pandemic. Thanks to all for their service.

1. SPSP Steering Committee

Chiara Ambrosio, Justin Biddle, Julia Bursten, Till Grüne-Yanoff, Catherine Kendig, Sabina Leonelli, Alan Love, Matthew Lund, Joseph Rouse, Eric Weber

2. Organising Committee of the 2020 Conference of SPSP

Chiara Ambrosio, Justin Biddle, Julia Bursten, Till Grüne-Yanoff, Sabina Leonelli, Alan Love, Matthew Lund, Joseph Rouse, Eric Weber (programme chair)

3. Local Organising Committee of SPSP/SRPoiSE 2020

Catherine Kendig, Kevin Elliott, Sean Valles (co-chairs)

Robert Pennock, Michael O'Rourke, Arthur Ward, Gregory Lusk, Isaac Record, Eleanor Louson, Heather Douglas, Ted Richards, and Robyn Blum

Additionally, the SPSP committee is grateful to the organising committee of the Consortium for Socially Relevant Philosophy of/in Science and Engineering (SRPoiSE) for its enthusiasm in working together with us to ensure co-hosting and overlap between the two conferences. We are sorry to lose this carefully orchestrated opportunity for exchange and hope it will be possible to re-organise it in the future.

III. SPSP Statement on the Covid-19 Pandemic

It is with great sorrow that we had to first postpone and then cancel the 8th biennial conference of the Society for Philosophy of Science in Practice (SPSP), which was originally scheduled for 7-10 July 2020. The timing of the COVID-19 outbreak, which became a pandemic at the same time as the programme for this conference was being finalised, made international travel to the conference impossible for the vast majority of delegates, including the steering committee and keynote speakers. Additionally, the local organizers at Michigan State University faced insurmountable practical obstacles, just like the organizers at any other institution would face under the current circumstances. For a conference scheduled in July 2021, several arrangements would have had to be made before the end of 2020. Rooms for sessions had to be booked and lodging for participants had to be reserved. This was impossible, because hosted events on campus were forbidden at MSU at least until early 2021. Another important element was that the numerous pre-pandemic options available in 2019 and early 2020 for local fundraising at MSU became much more restricted, despite valiant efforts by the local organising team in pre-Covid19 times. This meant that the registration fee would have had to be significantly higher than intended, at a time where the majority of research institutions are freezing their research and travel budget. This would have sharply curtailed the number of participants who would have been able to participate in an in-person SPSP conference. These and other problems were not specific to MSU. If another department in the US or a different country considered taking over for MSU, issues of public health safety, equity, and financial austerity measures or similar problems would have arisen. This meant that a real, on campus SPSP conference that was both safe and inclusive was not possible in 2021. We have considered an online conference, but have decided against it. An online meeting requires human and financial resources (e.g., for paying an online platform and hiring support staff to monitor the sessions) that SPSP as a society does not have at this moment, especially at such short notice and in a climate of serious financial restrictions. We wish to thank all our delegates for submitting wonderful abstracts, all the many people who helped us to set up this meeting, our webmaster Michel Durinx, and the excellent keynotes who had accepted to present at this conference: Till Grüne-Yanoff, Kristie Dotson, Sean A. Valles, Kylie Whyte and Karen Barad.

The steering committee will monitor the ongoing pandemic situation and discuss plans for the future over the coming months, with the goal of deciding when it is safer to schedule the next in-person meeting. In the meantime, we welcome feedback and informal expressions of interest for hosting the conference (please contact Sabina Leonelli at s.leonelli@exeter.ac.uk). Our amazing newsletter team will continue to produce informative magazines through which our community can read the latest news (including about SPSPers responses to the COVID-19 pandemic), while the mailing list will continue to serve as means to keep in touch with each other and support ongoing SPSP-related initiatives around the world. Keep well and until soon!

The SPSP Steering Committee and Local Organising Committee

IV. Abstracts of Symposia

In order to preserve the interactive and collaborative scholarship suggested in our accepted symposium proposals, symposia abstracts are presented by symposium rather than alphabetically by individual talk, and talks are listed in their intended order of presentation. Our 12 accepted symposia are listed alphabetically according to the surname of the first organizer.

1. Research on Organisms in Practice

Organizer: Rachel Ankeny

This symposium is based on an ongoing collaborative research project, “Organisms and Us,” which seeks to explore how researchers learn from and 'think with' non-human organisms in 20th and early 21st century science using a range of interdisciplinary methodologies and with particular focus on science in practice. We will present our findings to date with emphasis on diverse methodologies including philosophical and historical analysis, qualitative fieldwork, and quantitative data analysis. We explore a range of research activities associated with projects focused on biological organisms to investigate epistemological questions such as what make certain organisms serve as valid representations or as the basis of justifiable inferences and extrapolations; why do researchers choose particular organisms and how do their research practices affect their continued use over time; and what criteria or conditions serve as the basis for generalizations from organisms in different disciplines or fields. The symposium will be of interest not only to those working on the life sciences, but to philosophers of science in practice focused on modeling, explanation, contemporary team research practices, and field/discipline building, and especially for those who wish to understanding the broader contexts associated with science in practice that must be considered when seeking answers to epistemological questions.

1A: Organisms as Models between Representation and Intervention

Sabina Leonelli and Rachel Ankeny

There is no shortage of philosophical on the use of organisms as models, some of it drawing heavily on historical and sociological work which tends to focus on the standardisation of organisms as a critical stage in related research processes. This paper builds on this literature while taking up a new philosophical angle, namely the idea that the environment, experimental set-ups, and other conditions in which the organism is situated are critical to its use as a model organism (for an early discussion of this approach in historical context, see Ankeny et al. 2014; cf. Nelson 2017 on scaffolds). In such cases, the organism itself is only one component in a more complex model. Hence we explore how material systems can ground inferences, extrapolations, interventions and representations made using these organisms, using a series of examples based on recent historical cases including the use of cages and other housing systems with rodents and farm animals (Ramsden 2009; Kirk & Ramsden 2018), and various types of field experiments and stations (drawing on Kohler 2002 and de Bont 2014 among other historical work). We argue that this type of situatedness is a critical part of the repertoires centered on experimental organisms (Ankeny & Leonelli 2016), and is thus critical to explaining how organisms are used in scientific research practice. This analysis assists us in making sense of the variety of modelling activities associated with the use of organisms in laboratories and elsewhere across various fields within biology (e.g., Meunier 2012; Love & Travisano 2013; Germain 2014; Gross & Green 2017 among many others), and particularly in clarifying the link between the representational role of organisms and the extent to which they serve as tools for biological, biomedical and environmental intervention even in the absence of a sophisticated explanatory apparatus.

Ankeny and Leonelli. 2016. Repertoires: A Post-Kuhnian Perspective on Scientific Change and Collaborative Research. *Studies in History and Philosophy of Science Part A* 60.
Ankeny, Leonelli, Nelson, and Ramsden. 2014. Making Organisms Model Human Behavior: Situated Models in North-American Alcohol Research, since 1950. *Science in Context* 27 (3).
de Bont. 2014. *Stations in the Field*. Chicago University Press.
Germain. 2014. From Replica to Instruments: Animal Models in Biomedical Research. *History and Philosophy of the Life Sciences* 36 (1).
Gross and Green. 2017. The Sum of the Parts: Large-scale Modeling in Systems Biology. *Philosophy and Theory in Biology* 9 (10).
Kirk and Ramsden. 2018. Working across Species down on the Farm: Howard S. Liddell and the Development of Comparative Psychopathology, c.1923--1962. *History and Philosophy of the Life Sciences* 40 (1).
Kohler. 2002. *Landscapes and Labscapes*. Chicago University Press.
Love and Travisano. 2013. Microbes Modeling Ontogeny. *Biology & Philosophy* 28 (2).
Meunier. 2012. Stages in the Development of a Model Organism as a Platform for Mechanistic Models in Developmental Biology: Zebrafish, 1970–2000. *Studies in History and Philosophy of Science Part C* 43 (2).
Nelson. 2018. *Model Behaviour*. Chicago University Press.
Ramsden and Adams. Escaping the Laboratory: The Rodent Experiments of John B. Calhoun & their Cultural Influence. *Journal of Social History* 42 (2009).

1B: Synergistic Explanations for Choosing Embryos in a New Multivariable Analysis of Organism Choice

Nathan Crowe, Michael Dietrich, Sara Green and Sabina Leonelli

An often-used explanation for a researcher's choice of organism is to invoke the "Krogh Principle." Coined by Hans Krebs to describe Krogh's 1929 exposition of how he chose his research organisms, it stated that "[f]or such a large number of problems there will be some animal of choice, or a few such animals, on which it can be most conveniently studied" (Krogh, 1929). In general, Krogh's Principle has been shorthand for how arguments about convenience have overwhelming driven a biologists choice of experimental organism, particularly when researchers make arguments about the "best" organisms for particular problem. We find Krogh's Principal far too simplistic, and have created a list of twenty individual criteria that can be both nested and function synergistically to explain organism choice. Our criteria were devised from related philosophical interpretations of theory and model choice, a broad literature review of organism choice discussions, field work, and historical scholarship. The criteria provide a new lens for interpreting how researchers choose organisms, giving philosophers and historians of science concrete explanatory reference points to understand how these choices change across time and space. As an example of how these philosophical criteria can be used in historical analysis, we look at how Xenopus changed in popularity as a research organism in developmental biology between the 1950 and 1980s. Instead of assuming that Xenopus essentially outcompeted other organisms as the "best" one for developmental biology research, we show that the synergistic relationships between our criteria changed over these decades, which effectively reframed the definition of what a researcher considered a 'better' organism and drove changes in organism choice.

Dietrich, M., R.A. Ankeny, N. Crowe, S.Green, and S. Leonelli. How to Choose Your Research Organism. *Studies in History and Philosophy of Biological and Biomedical Sciences*, Online First, December 26, 2019.

1C. Fieldwork with Organisms: How to Do Practice-Informed Philosophy of Biology

Rachel Ankeny

A critical part of our broader research project "Organisms and Us: How living things help us to understand our world" is engaging in philosophically, historically, and ethnographically informed fieldwork to interview scientists and observe current practices in a variety of settings. This paper describes potential methods underlying such fieldwork approaches, drawing on relevant literatures from science and technology studies and history of contemporary science, including how they can be used to frame epistemological questions and explorations. I focus on my fieldwork in Australia with attention to three key topics that have emerged: the intersection between taxonomic and naming practices and regulation and policymaking with regard to native and non-native animals; disease modelling with non-standard organisms; and the use of genomic sequencing with native organisms to engage with global big science initiatives. I illustrate how fieldwork helped to articulate the key concepts and conflicts associated with each of these case studies, and outline ways in which reliance on published literature or more traditional approaches would have been likely to obscure many of the key issues. I also explore potential pitfalls associated with this type of fieldwork, including limitations due to choices of case studies, voices and perspectives that go unheard, and difficulties in generalizing from a discrete set of cases of practices, and how such considerations can be mitigated when planning fieldwork and developing philosophical arguments grounded within fieldwork.

1D. Measuring Research Biodiversity and Practices of Generalization in Postgenomic Biology

Michael Dietrich

When the NIH decided that a small set of organisms would be designated model organisms, some biologists feared that this designation would cause a drop in the diversity of organisms used in biological research. Such a change in research biodiversity could have a major impact on the kinds of questions posed, the methods used, the data collected, and the knowledge generated by postgenomic biology. Using publications from 1970 to the present in major disciplinary journals, we have created a record of organism use in genetics, evolutionary biology, ecology, cell biology, developmental biology, neurobiology, and physiology. Using two measures of biodiversity, we compare the diversity of research organisms used in each of these biological disciplines. Our expectation is that the designation of model organisms will have had little impact in fields such as ecology and evolution, but more significant impact in fields such as cell biology, which are more dependent on NIH funding. This is not to say that research biodiversity is driven by funding. Comparison of changes in research biodiversity across these different disciplines highlights the extent to which different areas of biology require organismal diversity for their research questions. It also reveals the extent to which some areas of biology are willing to assume, often without much justification, that research on a small number of samples from a single species is sufficient grounds for generalized conclusions. We will use this new data on research biodiversity to understand the different circumstances that inform generalizations from organismal results in different biological disciplines.

2. "You can Observe a lot by Watching": Imaging and Visualizing Individual Neurons

John Bickle, Ann Sophie Barwich & Lauren Ross

Our symposium title is a classic "Yogi Berra-ism," but it humorously captures a crucially important practice in current neuroscience: the use of imaging technologies to visualize structure and function in the central nervous system. While philosophical analyses of these technologies have overwhelmingly focused on their application to large-scale neural systems, the talks in this symposium explore their applications to single neurons. Our goal is to present some actual details of various single neuron imaging techniques and to show how an appreciation of their methodologies advances familiar discussions in current philosophy of science.

Single neuron techniques continue to be prominent in the mainstream fields of cellular and molecular neurobiology—fields which are the focus of the largest population of practicing neuroscientists. Philosophers' almost unanimous focus on larger-scale neuroimaging techniques leave out important lessons about "what we know about the brain," and about the practices by which we came to know this.

All three philosophers on our symposium engage with cellular and molecular neurobiologists. These talks will clarify how lower-level, molecular information about single neurons provides knowledge about neural systems. Ann Sophie Barwich describes a recent advance in microscopy that enables rapid live-stream imaging of small living organisms or intact brain tissues of larger animals. She argues that this tool, contrary to "dynamical" appearance, provides a new resource for a distinctively reductionistic approach in neuroscience that offers mechanistic explanations contingent on ongoing updates of molecular dynamics in neurons. She illustrates this claim by describing a yet-unpublished study on olfaction demonstrating that seemingly higher-level effects in mixture perception (i.e., suppression and enhancement) are not emergent properties of the system, but effects of elemental coding mechanisms at the periphery. John Bickle introduces techniques by which activities in individual neurons are imaged by their expression of immediate early genes. These imaging techniques have been especially useful to memory researchers investigating the molecular mechanisms by which individual neurons get allocated into specific memory traces, or engrams. Their use demonstrates how molecular pathways are studied "on the lab bench" and how their activities get linked experimentally to higher-level cognition, such as memory. Lauren Ross examines common "tracer" techniques in neuroscience, in which active transport processes in neurons are "tagged" in order to visualize the length, directionality, and spatial location of single neurons and neural connections. Her talk focuses on neuroscientists' use of the rabies virus as a visualizable tag that marks some material so that is can be followed as it moves through a causal process. Ross focuses on the particular types of causal information the use of such tracers do (and do not) provide. Despite superficial appearances, neuroscientific tracer methodologies do not support either mark transmission or standard "mechanist" accounts of causation. Instead, they are best understood within an interventionist framework as providing information about causal and molecular "pathways," notions which are poorly accommodated by the mechanist framework. These talks introduce state-of-the-art techniques from contemporary neuroscience-in-practice that will be novel to most philosophers, but state out positions from their uses that speak to familiar philosophical concerns.

2A. Imaging the Living Brain: Reductionism Past the Resting State
Ann-Sophie Barwich

Imagine studying an entire brain in action, live stream in all its complex interactive signaling, down to the single-cell level. Can we now look at the interrelated dynamics that fixed-tissue brain slices didn't cover, or where previous imaging techniques failed in resolution? Would this present the nail in the coffin for reductionist accounts in the life sciences? Contemporary advances towards a process philosophy of biology seem to suggest as much (Dupré 2012; Nicholson 2019). However, this talk demonstrates, the opposite is the case.

Advanced methods in neuroscience promote a reductionist approach to cognition in light of the detailed molecular and cellular mechanisms (Bickle 2015). An excellent example for this is a recent invention by the Hillman lab at Columbia (Bouchard et al. 2015): SCAPE (Swept, Confocally-Aligned Planar Excitation microscopy), a tool for 3-dimensional, rapid live-stream imaging of small living, freely moving organisms (C. elegans, drosophila) and living brain tissues of larger animals (mice).

Drawing on the results of a yet unpublished study in olfaction with SCAPE (Xu et al. 2019), I argue for a theoretical revolution embodied by such modern real-time molecular imagining tools. My paper explicates the causal difference between chemical and biological models of the olfactory stimulus, and how this difference points at the real mistake in previous reductionist explanations: we had operated with the wrong causal elements. To understand which chemical features are causally relevant, and in what order and combination, requires knowledge about the receptors concerning their binding behavior and mechanisms. Prior to SCAPE, the molecular dynamics of these mechanisms were inaccessible. What recent insights with SCAPE into the molecular mechanism of odor coding reveal is that the visible hallmarks of biological processes (with all their contextuality, unpredictability, irregularity, functional promiscuity) are fundamentally accountable down to their molecular and cellular causes. The appearance of emergence and irreducibility of these phenomena to essential material factors is a consequence of the misidentification of the principal elements and the misframing of the causal mechanisms responsible.

With the case of SCAPE, I will demonstrate that new imaging tools, paired with genetic techniques, allow for a targeted investigation of what previous technologies had required to study and therefore conceive of as different causal levels, such as the behavior of individual cells versus cell populations. The response of cells in the context of other cells as a population can now be modeled on one causal plane and integrated into a causal mechanistic explanation. Cellular mechanisms no longer provide mere details to supply higher-level computational models of physiological processes but constitute the material foundation from which to derive better neuroscientific theories. Modeling of dynamical systems in neuroscience cannot proceed without a revised and detailed conception of reductionism, which yields mechanistic explanations as contingent on ongoing updates of molecular dynamics.

Bickle, J. 2015. Marr and Reductionism. *Topics in Cognitive Sciences* 7.
Bouchard, M. B., V. Voleti, C. S. Mendes, C. Lacefield, W. B. Grueber, R. S. Mann, R. M. Bruno, and Hillman, E.M. C.. 2015. Swept Confocally-Aligned Planar Excitation (SCAPE) Microscopy for High-speed Volumetric Imaging of Behaving Organisms. *Nature Photonics* 9(2).

Dupré, J. 2012. *Processes of Life: Essays in the Philosophy of Biology*. Oxford University Press.
Nicholson, D J. 2019. Is the Cell Really a Machine? *Journal of Theoretical Biology* 477(21).
Xu, L., Wenze L., V. Voleti, Hillman, E.M.C., and Firestein, S. 2019. Widespread Receptor Driven Modulation in Peripheral Olfactory Coding. bioRxiv: 760330.

2B. Imaging Activity in Individual Neurons vis Immediate Early Gene Expression: Directly Linking Mind to Molecular Activities on the Lab Bench
John Bickle.

The 1980s and 1990s were peak years for discovery and development of new tools for imaging individual neurons and their activities in vivo. Many of these techniques are still in wide use today. Memory research was an especially prominent field for developing these single-cell functional imaging techniques; their use has provided a rich understanding of the molecular mechanisms of memory. Their use illuminates otherwise puzzling clams by neurobiologists to have "linked" mind to molecular activities. Expression of the immediate early gene (IEG) c-fos was first shown to be regulated in individual neurons in dynamic patterns, suggesting that its transcription and translation flag recent neuronal activity. This observation spurred the search for additional IEGs that could serve as markers of neuronal activity and for new methods to image the specific neurons expressing them. A decade after the c-fos discovery, expression of another IEG (activity-related cytoskeleton-related protein (Arc)) was confirmed as another marker of recent neuronal activity. This was confirmed using a then-newly developed imaging technique, cellular compartment analysis of temporal activity by fluorescent in situ hybridization (catFISH) (Guzowski et al. 1999). In these experiments, Arc expression in individual hippocampus neurons was "qualitatively and quantitatively similar" to the activation of hippocampus 'place' neurons, previously detected by single-neuron electrophysiology.

Other memory researchers found novel applications of these methods. Han et al. (2007) merged images of Arc-expressing neurons in the lateral nucleus of the rodent amygdala with images of green fluorescent protein-labeled neurons, which had been infected with a virus vector carrying the cloned transgene for a transcriptional activator, cyclic AMP-responsive element-binding protein (CREB). Infected neurons overexpressed CREB and were many times more likely to be included in the engram for a specific tone-fear memory than were uninfected neurons. This result suggests that increased CREB activation biases individual neurons to "win" in a competition for inclusion in a specific memory trace. Yiu et al. (2014) then used a variety of techniques for directly manipulating CREB activation levels in individual lateral amygdala neurons, including more recent optogenetic and chemogenetic approaches, and for imaging the most active neurons. Combined results clearly indicated that allocation of individual neurons into specific tone-fear association memories is at least partially biased by relative neuronal activity at the time of training. More active neurons at the time of training are more likely to be allocated to the specific memory trace for the association. These single-cell imaging techniques thus reveal the role of specific molecules, for example CREB, in higher-level cognitive functions.

Guzowski, J.F., McNauughton, B.L., Barnes, C.C., and Worley, P.F. (1999) Environment-specific Expression of the Immediate Early Gene Arc in Hippocampus Neuronal Ensembles. *Nature Neuroscience* 2(12).

Han, J.-H., Kushner, S.A., Yiu, A.P., Cole, C.J., Matynia, A., Brown, R.A., Neve, R.L., Gudowski, J.F., Silva, A.J., and Joseelyn, S.A. (2007) Neuronal Competition and Selection During Memory Formation. *Science* 316(5823).
Yiu, A.P., Mercaldo, V., Yan, C., Richards, B., Rashid, A.J., Hsiang, H.I., Pressey, J., Mahadevan, B., Tran, M.M., Kushner, S.A., Woodin, M.A., Frankland, P.W., and Josselyn, S.A. (2014) Neurons are Recruited to T Memory Trace Based on Relative Neuronal Excitability Immediately Before Training. *Neuron* 83(3).

2C. Tracer Methodology and Mark Transmission: Lessons from Neuroscience
Lauren Ross

This talk examines the use of tracer experiments in neuroscience to study neural connectivity in brain and nervous tissue. These experiments are found in a variety of scientific fields, where they involve “marking” some property with a visualizable tag so that this property can be followed as it moves throughout a causal process. For example, radioactive markers have been used to trace metabolic pathways in biochemistry, vascular pathways in physiology, and prey-predator food chains in ecology. These techniques bear striking similarity to mark transmission accounts of causation and this similarity has been noted by proponents of these accounts, including Reichenbach and Salmon (Reichenbach 1971; Salmon 1984). Despite this apparent similarity, little attention has been paid to how these techniques are actually used in scientific practice and whether this practice supports these philosophical views. The little attention these techniques have received in philosophy is also surprising given their importance in this scientific domain, where they are viewed as a “fundamental technique” and “the primary method for visualizing brain networks in all areas of neuroscience” (Arenkiel 2015). In this talk I analyze the rationale behind these tracer experiments and how they contribute to the philosophical literature on causation. I focus on the use of viral tracers–in particular the rabies virus–in elucidating sequences of neural connections. The rabies virus travels along sequences of neurons as a part of its natural infective capacity–after entering the body (typically through a bite wound) the virus infiltrates local neurons and moves from one neuron to the next, in a retrograde fashion, until it arrives in the central nervous system of the host. Scientists exploit the transsynaptic movement of this virus, by inserting in into particular neurons and watching it flow into the upstream neurons that they are connected to. I examine this case study in order to address (1) how these tracers are used to study neural systems, (2) what types of causal information they provide (and what types they do not provide), and (3) how an analysis of these tracers bears on the philosophical literature on causation. I argue that scientific tracer methodology does not support a mark transmission account of causation, but that these methods do involve a principled rationale and that they can provide information about causal systems. I indicate how the causal information they provide differs from standard “mechanistic” information (Craver 2007) and how this supports a more diverse picture of causal reasoning in neuroscience than in accepted by mainstream philosophical views.

Arenkiel, B. R. (2015) *Neural Tracing Methods*. Humana Press, Saskatoon, SK, Canada.
Craver, C. F. (2007) *Explaining the Brain*. Oxford University Press.
Reichenbach, H. (1971) *The Direction of Time*. University of California Press, Berkeley.
Salmon, W. C. (1984) *Scientific Explanation and the Causal Structure of the World*. pages 135–157. Princeton University Press, Princeton, New Jersey.

3. Levels of Organization and Biological Practice

Daniel S. Brooks, James DiFrisco, Markus I. Eronen, James Griesemer, Alan C. Love, Angela Potochnik, & William C. Wimsatt

A widespread idea in the biological sciences is that the natural world is separated into 'levels of organization.' This image in complete or abridged form figures centrally in many scientific and philosophical discussions, exhibiting roles as both major organizing principle and conceptual tool. Despite this, the precise nature and significance of 'levels' and its contributions to scientific and philosophical theorizing are rarely developed in detail (Love 2012), and instead often linger unchecked as auxiliary assumptions that do significant argumentative work (Griesemer 2005). Accordingly, this situation has provoked both skeptical attitudes towards (Potochnik and McGill 2012; Potochnik 2017, ch. 6; Eronen 2015), and in turn defenses of (DiFrisco 2017; Brooks 2017), the levels concept in recent philosophy of science literature.

The first part of this session addresses what we see as the continued promise of levels of organization for future scientific efforts and reasoning. For one thing, the levels concept continues to offer, at least in part, an ambitious ontological vision of the basic structure of the natural world. Though substantiating this vision has proven deeply challenging, we believe there is reason for optimism. We ground this optimism in the success that the concept exhibits in scientific practice.

In contrast to the optimism about levels of organization professed in the first half of this session, others contend that concepts of levels of organization are obscure, inaccurate, or otherwise problematic. The second half of this session will thus apply pressure to 'levels enthusiasts' in light of more skeptical treatments of the concept. Though both sides of this debate proceed from appeals to levels of organization in scientific practice, they draw opposed conclusions about the concept. The skeptics see appeals to levels of organization as having significant downsides, such as introducing confusion, suggesting implications that fail to obtain, and obscuring important findings that violate our assumptions about levels.

Following the individual presentations, we will convene a panel discussion with all contributors to address the prospects, promises, or pitfalls of 'levels' for scientific purposes in light of new and emerging perspectives on the concept (Brooks, DiFrisco, and Wimsatt in press). Importantly, both camps appear to agree that renewed attention to the levels concept is strongly needed to counteract the general tendency to treat the notion and its influences as primitive.

Brooks, Daniel S. 2017. In Defense of Levels: Layer Cakes and Guilt by Association. *Biological Theory*. 12(3).

Brooks, Daniel S., James DiFrisco, and William C. Wimsatt. in press. *Levels of Organization in the Biological Sciences*. Edited volume (15 contributed papers), MIT Press

DiFrisco, James. 2017. Time Scales and Levels of Organization. *Erkenntnis*, 82(4).

Eronen, Markus I. 2015. Levels of Organization: A Deflationary Account. *Biology and Philosophy*. 30(1).

Eronen, Markus I. 2019. The Levels problem in Psychopathology. *Psychological Medicine.*

Eronen, Markus I. in press. "Levels, Nests and Branches: Compositional Organization and Downward Causation in Biology." In Brooks, DiFrisco, & Wimsatt *Levels of Organization in the Biological Sciences*.

Griesemer, James R. 2005. The Informational Gene and the Substantial Body: On the Generalization of Evolutionary Theory by Abstraction, in *Idealization XII: Correcting the Model*, Martin R. Jones and Nancy Cartwright (eds.), Amsterdam: Brill, pp. 59–116.

Love, Alan C. 2012. Hierarchy, Causation and Explanation: Ubiquity, Locality and Pluralism. *Interface Focus* 2(1).
Potochnik, Angela. 2017. *Idealization and the Aims of Science*. University of Chicago Press.
Potochnik, Angela. in press. “Our World is Not Organized into Levels.” In Brooks, DiFrisco, & Wimsatt *Levels of Organization in the Biological Sciences*.
Potochnik, Angela and Brian McGill. 2012. The Limitations of Hierarchical Organization. *Philosophy of Science*, 79(1).

3A. A Fragmentary Account of Levels of Organization
Daniel S. Brooks

The concept of ‘levels of organization’ exhibits profound variation in its content, i.e. what it expresses, across different instances of its usage. This polysemic profile renders collecting a unified ontological account of the concept extraordinarily difficult due to competing definitional criteria that find shelter under the term’s multiple meanings. This in turn has motivated skeptical treatments of the concept, which argue (i) that scientific and philosophical work attributed to ‘levels’ can be streamlined by deflating the term into more precise notions, such as composition and scale or, more ambitiously, (ii) that the polysemic profile surrounding ‘levels’ is indicative of the concept’s principled inability to deliver lasting importance in its own right. Thus, the apparent ambiguity surrounding ‘levels’ appears to not only block progress on understanding what levels are, it also threatens the concept’s attributed centrality to scientific thinking in general. In this paper, I attend to these worries by developing an approach that situates ‘levels’ within an interest-relative matrix of operational usage within scientific practice. This approach proposes the notion of a fragmentary concept to account for the polysemy surrounding the levels concept in scientific contexts. Rather than speaking of a stable meaning in such concepts across different instances, it is more accurate, and productive, to speak of distinct and mutually combinable content fragments (CFs) of meaning (hence ‘fragmentary’ concept), which are set and weighted by scientific users in relation to their aims, values, and empirical setting. These CFs represent distinct core attributes of the levels concept whose contents are set by linking them to scientific tasks whose execution specify, in a piecemeal way, what the levels concept expresses in a given instance. In this way, ‘levels of organization’ is revealed to be highly customizable to a variety of usage circumstances consistent with widely distinct scientific contexts.

Viewing ‘levels’ as a fragmentary concept provides several benefits, and opportunities, to more adequately reconstruct the nature of ‘levels’ and its significance to scientific thinking. For one thing, this approach directly complements its unifying rationale of structuring scientific problems. Combining these two elements of its usage (customizable variation and unifying epistemic goal) reveals a ‘fragmentary account of levels’ whereby different instances of the concept can be ordered into lineages of usage, which in turn permit assessment of the concept’s contributions to achieving scientific goals. Coupled further with the concept’s overarching character as a scientific doctrine, the ontological promises of ‘levels’ can be substantiated as trends in regularity and predictability accumulated in scientific usage of the concept for certain goals over time. To illustrate this, I consider several examples where the operational framework of the fragmentary account delivers robust ontological content promised by the levels concept. I conclude with some general upshots for an account of the ontology of levels.

3B. Levels, Perspectives and Thickets: Toward an Ontology of Complex Scaffolded Living Systems
James Griesemer

Complex systems appear to be hierarchically organized and the idea of compositional levels of organization is one classic strategy to capture this type of complexity. However, compositional levels are inadequate for capturing the complexity of complexity. Wimsatt's ontology for complex systems — comprised of compositional levels, theoretical perspectives, and causal thickets — is mobilized in an argument that living developmental systems are typically scaffolded in ways that not only make them "interactionally" as well as descriptively complex, but with the effect that their system/environment boundaries change dynamically through the developmental process. This shift of perspectives toward starting with complex, thickety phenomenal experience out of which order is resolved, rather than with imagined elemental units out of which order emerges, points to an early step in scientific practices more generally: committing to a way of looking, to a mode of attention, for phenomena to track is a key step in initiating scientific work, regardless of whether the practice is empirical or theoretical. Contrasts between these "ways of working" are illustrated by examples of modeling and empirical practices in biology.

3C. Manipulating Levels of Organization
Alan Love

Despite their widespread invocation in scientific practice and pedagogy, levels of organization as hierarchical representations have come under increasing philosophical scrutiny. However, another significant dimension of scientific practice is manipulation. Scientists not only purport to represent levels, they also claim to manipulate them. In this paper I examine the manipulation of relatively stable configuration states in developmental biology with special attention to the origin of tissue-level organization from cell-level organization during embryogenesis. I use two forms of experimental practice for illustration: mixed cellular aggregates and tissue engineering. These successful experimental practices that rely on diverse forms of manipulation help to demonstrate that levels of organization cannot be reduced to principles of scale and do not support the claim that levels are only defined locally in terms of mechanisms. Additionally, they do not correspond to comprehensive, abstract "layer-cake" perspectives that are global in scope and map directly onto the disciplinary structure of the sciences. Overall, manipulation practices show that developmental biologists concentrate on transitions between particular levels, and the experimental establishment of these transitions contributes to interpreting cell and tissue configuration states as ontological levels of organization.

3D. Levels, Nests and Branches: Compositional Organization in Biology
Markus Eronen

The idea of hierarchical levels of organization is deeply rooted into contemporary biology and its philosophy. It refers to layers in nature, where entities at a higher level are composed of entities at the next lower level. Typical levels of organization are the molecular level, the level of cells, the level of organisms, and the level of populations. But are there really such levels in nature? In this talk, I will argue that when we take a closer look at scientific practice and the actual structure of biological systems, it turns out that hierarchical organization in nature is far more complex and messy than has been assumed. More specifically, due to the heterogeneity of biological components, actual compositional hierarchies in biology tend to be branching and tangled. This does not result in levels in the sense in which they have been traditionally understood. Consequently, the idea of levels should be treated as a heuristic idealization or abstraction that is only useful in some specific biological contexts. In most contexts, biological organization should be analyzed in terms of more well-defined concepts, such as scale and composition. I demonstrate the importance and usefulness of this approach by applying it to the debate on downward causation.

3E. Our World Isn't Organized into Levels
Angela Potochnik

Levels of organization and their use in science have received increased philosophical attention of late, including challenges to the well-foundedness or widespread usefulness of levels concepts. One kind of response to these challenges has been to advocate a more precise and specific levels concept that is coherent and useful. Another kind of response has been to argue that the levels concept should be taken as a heuristic, to embrace its ambiguity and the possibility of exceptions as acceptable consequences of its usefulness. In this talk, I suggest that each of these strategies faces its own attendant downsides, and that pursuit of both strategies (by different thinkers) compounds the difficulties. I will give special attention to responding to proposals by Alan Love, Daniel Brooks, and James Griesemer for the work accomplished by levels concepts. In contrast to their views, I believe the invocation of levels misleads scientific investigations more than it informs them, and this should be given fuller attention.

4. New Directions in Multiscale Modeling: Beyond Reduction and Emergence
Julia R.S. Bursten and Collin Rice

Multiscale modeling (hereafter MSM) has become a topic of significant recent interest, both among philosophers in the modeling literature and those working in inter-theory relations. MSM is widely used across the natural and social sciences, and it has become particularly prevalent across physics, biology, materials science, neuroscience, and economics in recent years. This symposium proposes to advance recent conversations on multiscale modeling by bringing together early-career philosophers working in diverse sciences. In contrast with much of the literature, which has focused on how to incorporate cases of multiscale modeling into existing philosophical discussions concerning reduction/emergence, representation, or explanation, this symposium considers practices of multiscale modeling on their own terms.

We aim to generate, through five case-driven talks, a comparative account of the varieties of MSM found across astrophysics, ecology, economics, nanoscience, and neuroscience. Through this aim, we hope to expand the kinds of conversation philosophers are having about MSM, with an eye to what makes particular MSM strategies successful and what sorts of explanatory and inferential capabilities different types of MSM approaches afford to users. For example, our symposiasts address questions such as: What justifies the use of a particular MSM strategy?, What practical constraints limit the construction of multiscale models?, and What is the connection between different modeling goals, scales, and the selection of MSM strategies?

The first two talks introduce new examples in physics. Bursten begins with a discussion of nanoscale alloys. She uses the case to illustrate how typical MSM modeling of alloys relies on length-scale separation between component models, and she discusses how modeling practices and scientific concepts respond when scale separation fails. Second, Jacquart examines MSM in astrophysics by examining the practice of modeling galaxy structure formation. She uses this case to demonstrate how MSM practice drives development, despite its scale challenges. Third, Jhun uses two case studies in economic policy analysis to turn attention to the use of multiscale models in constructing narratives that guide policy-making. Despite the ad-hoc ways in which models can inform policy, economists successfully use them to construct coherent stories about the possible effects of particular actions. Fourth, Haueis also examines purposes to which multiscale models are put. Using case studies in neuroscience, he introduces an account of MSM for exploratory and descriptive scientific practices, rather than explanatory ones. Fifth, Rice discusses the use of homogenization techniques to model heterogeneous landscapes in spatial ecology to advocate for a pluralistic approach to MSM strategies that goes beyond cases of scale separation.

Shared themes across these talks include (1) interest in the different ways in which component models can combine, especially when scale separation is challenged, (2) how these differences constrain the kinds of inferences available from the use of multiscale model, and (3) the extent to which MSM strategies are constrained or guided by relations between scales, modeling goals, or other practical constraints on scientific practice. Through these themes, we hope to expand philosophical conversations about MSM.

4A. Multiscale Modeling of Nanoscale Alloys: When Scale Separation Fails
Julia R.S. Bursten

In multiscale modeling of physical systems, dynamical models of higher-scale and lower-scale behavior are often developed independently and stitched together with connective or coupling algorithms, sometimes referred to as "handshakes." This can only be accomplished by first separating modeled behaviors into bulk behaviors and surface or interfacial behaviors. This strategy is known as "scale separation," and it requires physical behaviors at multiple length, time, or energy scales to be treated as autonomous from one another.

In this talk, I examine what makes this strategy effective — and what happens when it breaks down. The nanoscale poses challenges to scale separation, because at the nanoscale the physics of the bulk occurs at the same length scale as the physics of the surface. Common scale-separation techniques, e.g. modeling surfaces as boundary conditions, fail. Using multiscale modeling techniques to model the scale-dependent physics of nanoscale materials presents a new challenge whose solution requires conceptual engineering and new modeling infrastructure. These considerations suggest a view of physical modeling that is centered not around idealization or representation but around the interaction of component models in a multiscale model. This view of modeling generates a shift in the sorts of philosophical considerations due to theorizing about models.

To illustrate this view of modeling, I consider contemporary cases of multiscale models of nanoscale multi-metallic compounds. At macroscopic scales, most multi-metallic compounds are alloys, and alloy modeling is a common venue for multiscale modeling strategies in materials modeling. At the nanoscale, the relations among the metallic components of multi-metallic compounds shift, and with these shifts come changes in modeling approaches. I use these shifts to show how multiscale modeling strategies adapt in the face of collapsing distinctions between mesoscopic and microscopic length scales.

4B. Multiscale Modeling of Astrophysical Galaxy Formations

Melissa Jacquart

As computational power has increased, cosmologists have produced complex simulations of galaxy structure formation, ranging from individual galaxies, to galaxy cluster interactions, to the structural history of the entire universe. These models and simulations play critical roles in the reasoning process in astrophysics. However, their use also presents a host of challenges. Galaxy structure modeling takes place within the same physical theory, namely, the Standard Cosmological Model (which is the Lambda Cold Dark Matter (LCDM) parametrization of the Big Bang cosmological model). Since the Standard Cosmological Model is expected to hold throughout the entire universe, LCDM multiscale modeling practices are often criticized for their scale dependence, and failure to successfully span all scales: from individual galaxies through to cosmological scales of the entire universe. For instance, Massimi (2018) discusses the challenges LCDM faces going down from large scale structure formation to the meso-scale of individual galaxies. She calls this the "downscaling problem": while LCDM galaxy structure formation simulations successfully explain large scale structures of the entire universe, it faces challenges explaining the observed Baryonic Tully-Fisher relation at the level of individual galaxies.

In this talk I will argue that multiscale modeling (MSM) practices of galaxy structure formation should not indicate a problem about the explanatory power of LCDM. Rather, MSM is an effective strategy deployed in the face of complexity as it relates to galaxy structure formation. On the meso-scale, individual galaxy simulations must incorporate gas density, dark matter, and stellar feedback during the course of the simulation. The complex behavior of plasma underlies many astrophysical processes. The micro-scale processes of plasma affect the appearance of meso-scale astrophysical systems such as supernova remnants or accretion disks. These then drive stellar feedback, which largely influences the development and evolution of a galaxy. The multiscale model's ability to model a single galaxy well relies on the component model's ability to model various micro-scale physical properties well. Additionally, in order to macro-scale cosmological structures (hundreds of thousands of galaxies), an accurate account of individual and cluster galaxies is needed.

I first provide my preliminary analysis of the use of MSM in this context of galaxy structure formation simulations. I discuss how the relationship between these two dramatically different scales of modeling inform each other and drive each level's further development of simulations. I suggest that it is the MSM strategy is in fact developing knowledge, despite downscaling problem or tyranny of scales. The more fruitful understanding of modeling in astrophysical contexts, given the sheer complexity at play, is that the component models yield domain-specific information. Yet when taken in aggregate, they provide capacity for inferences. As such, success and inferential power of galaxy structure formation models rely on MSM techniques. What MSM practice does in the context of galaxy structure formulations is identify the points at which the modelling assumptions (namely the idealizations, representations, and levels of complexity) must be attended to, and where scale-specific models should not be expected to extend beyond their original domain.

4C. Multiscale Reasoning in Economics

Jennifer Jhun

There is a substantial amount of discussion in the philosophical literature on multi-scale modeling in physical and material contexts. But there is less such discussion when it comes to the social sciences. This paper earmarks economics as a promising area for multi-scale exploration for a number of reasons. First, there's a sense in which things that go on at one scale on their own do not reduce to what goes on at another, though goings-on at one scale may contribute to goings-on at another. Formal results such as the Sonnenschein-Mantel- Debreu theorem and the questionable success of the microfoundations program provide at least two prima facie reasons we should be suspicious of reductionist attempts in economics. The next natural step would be to consider how it is that economists use multiple models together. Second, economists in actual practice, such as those at central banks, often really do use multiple models in order to achieve monetary policy aims.

The literature is now over trod with commentary about the representational capacity of these (often idealized) models, so much so that not only is the discussion often quite divorced from economic practice, but it also makes the possibility of a realist interpretation of economics look rather tenuous. This paper shifts its attention to the role of models not just as representational devices, but also as devices for the construction of performative narratives. Models, when deployed in policy analysis, are part of the effort to construct coherent narratives that help guide action. This is an enterprise that often requires different parties to coordinate their expertise and resources. This role requires that models take on both the role of being explanatory - being able to be the kind of tool that people can use to ask what-if-things-had-been-different or why-questions — and performative — being the kind of tool that economists can use to effectively enact changes in the world. Doing so will highlight a number of features about economic models in particular that has escaped much of mainstream philosophical attention — for instance, why it might be that economic models furnish how-actually (as opposed to merely how-possibly) explanation.

The multiscale modeling framework, I propose, is one promising way of grounding this conception of model usage. To this end, I examine two case studies. One is explicitly multi- scale: integrated computable general equilibrium and microsimulation modeling strategies, with particular attention to income distribution effects. The second is more implicit: I ex- amine the process by which the U.S. Federal Reserve, via the construction of the Greenbook (now Tealbook) conditional forecasts which are then distributed at the Federal Open Market Committee meetings for discussion, offers policy recommendations. Finally, I suggest that these considerations actually point to a kind of pragmatic realism as the right attitude to have towards many economic models.

4D. Multiscale Descriptive Modeling in Data-driven Neuroscience

Philipp Haueis

In this paper, I argue that multiscale modeling can be used to explore systems whose organization at different scales is ill-understood, and to describe the relationship between those scales. My argument combines the claim that exploration and description should stand alongside explanation and prediction as a core function of models in scientific practice (Gelfert 2018) and the claim that descriptive models in neuroanatomy serve as standards based on which multiple different explanations of brain functions can be developed (Ankeny 2000, Haueis and Slaby 2017). Using a case study from data-driven neuroscience, I argue that descriptive multiscale models differ from multiscale explanatory models in the way they (1) select information to describe a system at a particular scale and (2) how they link such scale-specific descriptions to each other.

In explanatory modeling, scale-specific information is selected based on its explanatory relevance to a particular macroscale behavior of the system. In physics, for instance, mesoscale information about dislocation pile-up is selected because it is relevant to explain when a steel bar cracks (Wilson 2017). In contrast, descriptive models in neuroscience contain information that is relevant to many different behaviors. Descriptively relevant information is selected based on a fundamental presupposition about the system (Ankeny 2000, Haueis and Slaby 2017). In the case of multiscale modeling of cortical gradients—continuous progressions of functional and/or structural features across the entire cortex—the fundamental presupposition is that an anatomical hierarchy of feedforward and feedback connections underlies a signal processing hierarchy of input-output relations (Markov et al. 2014). Hierarchical cortical gradients can be modeled at the macroscale of whole-brain networks (Margulies et al. 2016), the mesoscale of cortical circuits (Hilgetag et al. 2019) or the microscale of individual cells and their dendritic branches and axonal connections (Paquola et al. 2019). The information in such multiscale gradient models is relevant to many different brain functions, e.g. how the brain processes visual information, allocates attention or integrates multimodal information.

To link scales, explanatory models use upper-scale boundary conditions to restrict the lower-scale model domain to regions where explanatorily relevant information is expected to occur, and lower-scale information is homogenized to be incorporated in the explanation of a macroscale behavior (Batterman 2013). In descriptive models, however, the explanatory relevance of lower-scale information is unknown and needs to be explored across the entire domain. Cortical gradient models, for instance, use macroscale functional connectivity patterns to parametrize a mesoscale model of cortical circuits across the entire brain (Wang et al. 2019, Demirtaş et al. 2019). Since the functional relevance of the macroscale patterns is itself only partially understood, these models use lower-scale information to determine what features the upper-scale data pattern could refer to. For example: macroscale gradients could both refer to gradual transitions in dendrite density at the microscale or changes in interlaminar inhibition at the mesoscale. By explicating parametrization and reference determination as distinct strategies of selecting and linking scale-specific information, my analysis reveals that besides explanation, and exploration and description are distinct functions of multiscale modeling in scientific
practice.

Ankeny, R.A., 2000. Fashioning Descriptive Models in Biology: Of Worms and Wiring Diagrams. *Philosophy of Science* 67.

Batterman, R., 2013. The Tyranny of Scales, in: Batterman, R. (Ed.), *The Oxford Handbook of Philosophy of Physics*. Oxford University Press, pp. 255–286.

Demirtaş, M. et al.., 2019. Hierarchical Heterogeneity across Human Cortex Shapes Large-Scale Neural Dynamics. *Neuron* 101 (6).

Gelfert, A. (2018). Models in Search of Targets: Exploratory Modelling and the Case of Turing Patterns. In Christian, A. et al.. (eds.): *Philosophy of Science. Between the Natural Sciences, the Social Sciences, and the Humanities*. Springer.

Haueis, P. and Slaby, J. (2017). Connectomes as Constitutively Epistemic Objects: Critical Perspectives on Modeling in Current Neuroanatomy. *Progress in Brain Research* 233.

Hilgetag, C.C., Beul, S.F., van Albada, S.J., Goulas, A., 2019. An Architectonic Type Principle Integrates Macroscopic Cortico-cortical Connections with Intrinsic Cortical Circuits of the Primate Brain. *Network Neuroscience* 3 (4).

Margulies, D.S. et al. (2016). Situating the Default-mode Network along a Principal Gradient of Macroscale Cortical Organization. *PNAS* 113(44).

Markov, N.T. et al. (2014). Anatomy of Hierarchy: Feedforward and Feedback Pathways in Macaque Visual Cortex. *The Journal of Comparative Neurology* 522(1).
Paquola, C. (2019). Microstructural and Functional Gradients are Increasingly Dissociated in Transmodal Cortices. *PLoS Biology* 17 (5).
Wang, P., Kong, R., Kong, X., Liégeois, R., Orban, C., Deco, G., van den Heuvel, M. and Yeo, T. (2019). Inversion of a Large-Scale Circuit Model reveals a Cortical Hierarchy in the Dynamic Resting Human Brain. *Science Advances* 5(eeaat7854).
Wilson, M. (2017). *Physics Avoidance*. Oxford University Press.

4E. The Practical Constraints of Multiscale Modeling
Collin Rice

Philosophical discussions of multiscale phenomena have tended to focus on their implications for long-standing debates about reduction, emergence, and relationships between fields of science. However, even if we set these metaphysical questions aside, a number of difficult practical modeling problems remain. In particular, in this paper, I use a number of examples from biological multiscale modeling to argue that the primary challenge facing these modelers is not how to metaphysically interpret their models, but is instead using idealization to bring their existing modeling techniques to bear on multiscale phenomena. The `best-case scenario' for multiscale modeling is when the dominant features of the system can be separated into distinct scales. When this occurs, scientists can effectively model the phenomenon by using modeling techniques designed for those particular scales (and type of processes). However, as Bill Wimsatt notes, ``In biology and the social sciences, there is an obvious plurality of large-, small-, and middle-range theories and models that overlap in unclear ways and usually partially supplement and partially contradict one another in explanations of interactions of phenomena at a number of levels of description and organization.'' (Wimsatt 2007, 179--80). In other words, in many cases of multiscale modeling in the biological and social sciences, such a clean separation of scales (or hierarchical organization) is not possible.

Therefore, in this paper I focus on the kinds of techniques multiscale modelers employ when the relevant features and processes of the system do not separate into distinct scales. In particular, I present cases in which biological modelers use homogenization techniques for modeling heterogeneous landscapes in spatial ecology and minimal modeling techniques in studying instances of self-organized criticality. As these cases make clear, rather than attempting to accurately model the complex inter-scale interactions of multiscale systems, scientific modelers routinely use a wide range of idealized modeling techniques designed to reduce complexity, identify stable patterns across causally heterogeneous systems, and make use of models developed for similar phenomena in other fields. These cases also show how the use of multiscale modeling techniques is often aimed at avoiding physically relevant details in order to draw inferences from very simple models that encode information from a variety of scales of the system. In addition to revealing several of the practical modeling constraints involved in the practice of multiscale modeling, these cases also illustrate a kind of explanatory independence that is largely missed by philosophical discussions of multiscale modeling that focus on reduction and emergence. Specifically, these cases show that many multiscale phenomena are unable to be explained or understood by any single kind of modeling strategy. As a result, rather than attempting to identify a single best strategy for modeling multiscale phenomena---e.g. top-down, bottom-up, or middle-out approaches---in most cases, multiple modeling techniques will be required to uncover the modal information required for accomplishing various modeling goals.

5. Putting Mechanistic Ideas to Work in the Practice of Medical Science

Lindley Darden

Medical researchers often search for and make use of knowledge about mechanisms: how they work, how they go wrong, and how to intervene on them. Philosophers are investigating the ways in which medical researchers apply mechanistic ideas to these tasks and in doing so enhance our understanding of this practice and can contribute to improving it. The kinds of mechanism of interest to medical researchers can be distinguished along various dimensions, such as the following:

1. Useful for diagnostics and drug development
   a. "Normal/healthy" versus disease mechanism
   b. Mechanism in the patient versus in an infectious agent
   c. General versus patient-specific mechanism
2. Useful for intervention and policy about a drug/therapy
   a. Mechanism of action of in a sample (human or animal) versus the target population (which could be one patient)
   b. Intended molecular or cellular mechanism of action verses side-effects
   c. Direct effects versus downstream effects

The symposium participants will address practical issues that arise in medical research on mechanisms distinguished across these dimensions. Darden and her scientific collaborators have developed a diagrammatic web-based framework for representing genetic disease mechanisms that is designed to aid medical researchers in organizing the burgeoning information from databases and the medical literature, as well as in finding targets for biomarkers and drugs. Bechtel will also direct attention to methods for finding disease mechanisms by analyzing the practices of network biologists who utilize large online databases to identify potential mechanisms in which different altered genes or proteins are components. By linking diseases not to individual genes but to underlying cell mechanisms, these researchers seek to reduce the heterogeneity and offer manageable explanations of disease states.

Rather than focusing on general disease mechanisms, Erasmus and Lean look at the practices of medical researchers in the second category above, namely work on assessing the mechanisms of actions of a drug, at both the populational and molecular levels. Erasmus provides an account of mechanistic evidence needed for predicting medical effectiveness. He provides, as others have not explicitly done, a detailed account of the criteria needed to reliably relate a drug's mechanism of action in the sample population to its mechanism of action in the target population. Lean focuses on the mechanism of action of drugs at the molecular level; in particular, the under-appreciated notion of specificity as it applies to molecular mechanisms and interventions on those mechanisms. He argues that scientists' understanding of how a medically relevant mechanism works is intimately connected with the means by which scientists aim to intervene on those mechanisms, which constitutes novel support for the dependence of mechanistic explanation on scientific practice.

Taken together, these papers offer a collection of new perspectives on the mutual relationships between philosophy and medical science in our understanding of the facets and contours of biological mechanisms.

5A. Using the New Mechanistic Philosophy of Science to Represent Genetic Disease Mechanisms
Lindley Darden, Kunal Kundu, Lipika R. Pal, Isaac Bronson and John Moult

Biologists and medical researchers often search for mechanisms. There is no standard way of representing disease mechanisms. This paper will discuss an ongoing collaboration between a philosopher of biology and computational biologists to use insights from the new mechanistic philosophy of science to tackle this problem. The typical components of mechanisms—entities and activities—are adapted to represent aberrant components—called "substate perturbations (SSPs)" and "mechanism modules (MMs)" Stages of genetic disease mechanisms begin with altered genetic material (DNA or chromosome aberration) and go through stages to the disease phenotype. There are types of perturbations at each stage, such as a single base change in DNA or an altered protein conformation, as well as the kinds of activities (or groups of entities and activities) that drive each change from stage to stage, such as altered protein folding. The result is an abstract, robust, and informative disease mechanism representation method, adapted from work in the new mechanistic philosophy of science to apply to genetic disease mechanisms.

Diagrams represent general abstract genetic disease mechanism schemas. They consist of stages of SSPs and MMs, from gene to phenotype. Using standard biological ontology terms, the successive stages are labeled: DNA to RNA to protein to protein-complexes to cell organelles to cells to tissues to organs to phenotype. The schemas include black boxes to indicate where additional research is needed. The diagrams indicate sites in the mechanism for possible drug targets and useful biomarkers.

This diagrammatic representation framework has been implemented in the web-based MecCog system (www.MecCog.org), which aids medical researchers in organizing knowledge scattered in the medical (PubMed) and textbook literature in easy-to-comprehend diagrams. A collaborative team of a philosopher, a computational biologist, and an undergraduate premed student built a schema representing the mechanism producing cystic fibrosis (CF). A medical expert in CF reviewed the schema and suggested improvements. Additional hypothesized schemas that represent complex trait diseases and cancer await review by experts. Future work is planned on Alzheimer's disease.

The MecCog system thus allows a perspicuous representation of what is known and, importantly, what is not known about general genetic disease mechanisms for specific diseases and specific gene mutations. It thus provides mechanistic explanation and understanding. It provides a means of collating evidence for and against each mechanism component and representing confidence scores, based on the researcher's judgment about strength of the evidence. It indicates uncertainty (low confidence score), ignorance (black box), and ambiguity (alternative branches). It indicates where future experiments should be directed and aids the identification of possible new drug targets and biomarkers. This work is an instance of fruitful interfield relations between philosophy of science and medicine.

5B. Practices of Network Biologists in Medicine: Discovering Mechanisms Altered in Disease
William Bechtel

In recent decades medical researchers have been compiling massive data—genetic, proteomic, etc.—based on diseased patients. In the case of cancer, for example, initiatives such as the Cancer Genome Atlas (TCGA) and the Catalogue of Somatic Mutations in Cancer (COSMIC) generated large online databases of genes mutated or exhibiting aberrant copy number in tumors. However, rather than directly pointing to a few altered genes as potential causes of tumors, these efforts revealed tremendous heterogeneity in the genes altered. Nonetheless, the databases have provided the resources for network biologists who have tried not just to use networks to represent the data but also to analyze networks to identify potential mechanisms in which different altered genes or proteins are components. By linking diseases not to individual genes but to underlying cell mechanisms, these researchers seek to reduce the heterogeneity and offer manageable explanations of disease states. My focus will be on articulating the practices of these network biologists. I will begin with the basic strategy of identifying clusters of highly connected nodes, treating them as mechanisms, and assigning activities to them by annotating nodes with information from Gene Ontology (GO). I will briefly describe one particular approach that diffuses activations from nodes for multiple altered genes across a network to identify mechanisms whose altered activity is inferred to explain the disease phenotype.

I will then concentrate on practices involving a novel strategy for using deep learning neural networks to analyzing disease phenotypes and efficacy of drug interventions. Although standard deep learning neural networks are powerful learners, able to learn and predict relations between inputs (representing, e.g., genes or drugs) and outcomes (representing phenotypical outcomes such as disease state or drug response) that generalize well to new cases, they are black boxes. This is because they connect each unit in one layer with every unit in adjacent layers and the pattern of connections is generally uninterpretable. Recently Trey Ideker's laboratory has adapted deep learning to create what they call visible networks—ones whose nodes and edges can be interpreted biologically. They do this by drawing upon bioontologies (GO and their own NeXO) to structure the neural network. Hidden nodes in their networks stand for terms in the ontology; nodes are connected only when the terms are linked in the ontology. Ontologies provide a compendium of hierarchically organized mechanistic knowledge about cells. When nodes representing terms in the ontology are altered as the network responds to different inputs and generates different outputs, researchers can interpret these nodes in terms of known biological mechanisms. These networks, DCell and DrugCell, perform slightly less well than fully-connected deep learning networks, but they enable researchers to construct testable hypotheses about the mechanisms altered in different states. This research program is still in its infancy so I will be focusing on emerging practices in this research program.

5C. Mechanisms for Predicting the Effectiveness of Medical Interventions
Adrian Erasmus

It has been argued that establishing causal claims in medicine should rely on both statistical and mechanistic evidence. These arguments are often framed as calls for augmenting evidence of correlation with evidence of mechanisms. While I agree with the spirit of this proposal, there is little in terms of detailed accounts of mechanistic evidence for particular aims of medicine beyond illustrations of individual cases where such evidence has helped justify causal claims. An important practical issue involves the role of mechanistic evidence in establishing claims about the effectiveness of medical interventions, especially when these claims are predictive. In this paper, I give an account of mechanistic evidence for making reliable predictions about medical effectiveness.

I begin by defining the kinds of mechanisms medical researchers typically aim to establish evidence of when assessing a drug's effectiveness – the mechanism of action of a drug. This notion refers to the specific interactions by which a biochemical substance produces its pharmacological effects. The mechanism of action of a drug typically references the molecular targets to which the drug binds and the specific actions that occur there. Take trastuzumab for breast cancer, which targets the HER2 protein, by binding to its extracellular juxtamembrane domain, thereby inhibiting proliferation and survival of HER2-dependent tumors.

While finding evidence of such mechanisms is important in establishing how the intervention in question produces its purported outcome, I argue that predicting a particular medical intervention's effectiveness in a target population, or for a target individual, requires more than just evidence of the mechanism of action of a drug. Some have argued that the reason we need mechanistic evidence is to mitigate possible confounding in clinical trials. However, this assumes that just having mechanistic evidence is enough. I suggest that a drug's mechanism of action can, in a sense, differ across populations due to several factors. I draw a distinction between a drug's mechanism of action in the sample population and the drug's mechanism of action in the target population, define each, and argue that the latter is largely ignored in discussions about mechanistic evidence in medicine. Moreover, I argue that neglecting the latter can limit our ability to make reliable predictions about its effectiveness.

I then provide an account of mechanistic evidence for predicting medical effectiveness. Improving such predictions requires evidence that the causal relationship in the sample population is sufficiently similar to that of the target population. This means we should aim for mechanistic evidence for the target population which satisfies conditions of stability (the evidence should indicate the extent to which the effect will be influenced by background conditions), specificity¬ (the evidence should relate directly to the mechanisms of action in question), and relevance (the evidence should include all and only pertinent information about the mechanisms in question). In doing so, I provide a way of assessing part of the evidence used in making predictions about the medical effectiveness.

5D. The Unavoidable Role of Practice in The Constitution of Molecular Mechanisms

Oliver Lean

Mechanistic models always involve abstraction—filtering out certain details about their objects in order to focus on others. Everyone agrees that abstraction is legitimate for various practical reasons: for computational efficiency; to make the models comprehensible to human intellects; to emphasise certain factors that matter in a particular context; or because those details are unknown. Nevertheless, Craver and Kaplan (2018) maintain that behind these merely pragmatic filters there exists, in reality, a complete set of explanatory relevant facts constituting a given phenomenon—something our abstract models have abstracted from. Crucially, rather than being just an exhaustive list of details, no matter how trivial, completeness only requires the "constitutively relevant" facts. Importantly, they hold that this constitutive relevance is ontic—a matter of how things really are, independently of our practical or epistemic concerns.

At the heart of this debate, evidently, are questions about explanatory relevance in constitutive mechanistic explanations: Are the facts that jointly constitute a mechanistic explanation for a phenomenon entirely a matter of our pragmatic interests and epistemic limitations? Or, as Craver and Kaplan argue, is there a way of understanding explanatory relevance that is independent of these?

In this paper, I introduce a new challenge to the idea of an ontically complete mechanistic explanation—one that comes to light when we look at scientific practice. My discussion will focus on contemporary practices of drug design—the field concerned with discovering, improving, and creating drug-based interventions on biological phenomena for both therapeutic and experimental purposes. It is important to observe that the action of a drug—the intervention mechanism, as I call it—is distinct from the target mechanism being intervened on. However, since the two of course interact, facts about the target mechanism are highly relevant for what makes an effective intervention mechanism. What makes a good drug depends, for example, on facts about the structure of the direct target, the mechanism of interaction with the candidate drug, and possible mechanisms of decoy interactions that are to be avoided.

These observations make the idea of ontic explanatory relevance harder to sustain. Craver and Kaplan adopt an interventionist notion of causality in clarifying relevance; however, their focus is on idealized interventions as discussed by Woodward, which (for good practical reasons) ignore how the intervention is to be performed. However, drug design's aim is to create actual interventions, whose interaction with their target is far more complex than their idealized counterparts. Crucially, then, pragmatic aims related to intervention do not just abstract details from an ontically complete mechanism, "filtering out" facts from the world according to our purposes and limitations. Instead, practical aims in manipulating biological targets can also add relevant details that were left out of a more idealized, purpose-neutral picture. This yields a general lesson: It is harder than supposed to portray scientific practice as merely abstracting details from a complete, ontic reality. Instead, the purposes, limitations, and abilities of scientific investigation are tightly intertwined with, perhaps even inseparable from, questions of the constitutively relevant facts in mechanistic explanations.

6. What is the Practice of Medicine?

Jonathan Fuller

Medicine is an interesting case for philosophy of science in practice because it is contested whether medicine is a science, a practice or both. The question 'what is medicine?' is an underexplored one, but should be seen as a field-defining question for the philosophy of medicine. Developing a general account of medicine, or even contemporary scientific medicine, is no trivial task. First, there is an incredible diversity of medical practices over history and across cultures, which makes it genuinely puzzling what these various practices all have in common. Second, demarcating scientific medicine may require demarcating science, which has proven vexatious. Third, it is unclear whether modern medicine is itself a science in addition to a practice. Finally, there are several rival interpretations of 'precision medicine', which is touted as the new dominant model of medical research and practice.

In a recent book titled *Philosophy of Medicine* (2019), Alex Broadbent throws down the gauntlet and asks: what is medicine? One way or another, the participants in this symposium seek to answer that challenge, providing responses from philosophical, clinical and integrated HPS perspectives.

First, Broadbent argues that to understand the practice of medicine we need a grip on its goal and core business. Broadbent argues that medicine's primary goal is to cure disease and its core business is understanding and predicting disease – medicine is fundamentally an inquiry. These are the elements that make diverse medical traditions recognizably medical, from African medicine to alternative medicine to modern mainstream medicine.

However, Chadwin Harris argues that Broadbent's inquiry model lacks a convincing account of the link between inquiry and cure. Harris divides medicine's core competencies into outward-looking (or patient-centric) and inward-looking (or profession-centric) competencies. The inquiry model correctly describes medicine's inward-looking competencies, but it falls short in describing medicine's outward-looking competencies.

Olaf Dammann objects to Broadbent's thesis that medicine is an inquiry, if by inquiry we mean something akin to science. There is an important distinction between medical science and medical practice. While medical practice is scientific, it is not itself a science. Medical practice does not generate generalizable knowledge but knowledge only of cases, of particulars. Thinking of medical practice as a science risks turning back the clock to a time before modern scientific medicine.

Next, Jonathan Fuller asks: what is modern medicine? Several historians argue that medicine became modern by the turn of the Twentieth Century with the rise of scientific medicine. Fuller argues that what made medical practice newly scientific (vs. pseudoscientific) was its engagement with forms of inquiry that sought and responded to evidence; this feature still demarcates modern scientific medicine from other contemporary medical traditions.

Finally, Zinhle Mncube asks, what is precision medicine? Mncube notes the various interpretations of precision medicine in healthcare, yet argues that there is a core concept at stake against which various interpretations can be judged. For instance, common pathobiological conceptions of precision medicine are impoverished because they are gene-centric and fail to capture the diversity of influences on health and disease.

6A. The Inquiry Model of Medicine

Alex Broadbent

In this paper I contrast two models of medicine, the Curative Model and the Inquiry Model, and favour the latter.

Medical traditions vary as widely as one can imagine in different times and places. Given the apparent variety of actual practices that are either claimed to be or regarded as medical, what, if anything, do they share—other than the fact we are inclined to refer all these wildly different practices as (at least partly) medical. Any adequate philosophy of medicine must respond to this question, and say something about what medicine is.

I consider a simple Curative Model, stating that the goal and business of medicine is to heal the sick. I distinguish the goal of medicine from its core business, that is, the exercise of some competence or skill that is characteristically medical and that is more than mere well-meaning but inexpert assistance, like bringing a sick relative a nice cup of tea and a pair of slippers. I defend the Curative Model's implication that cure is the goal of medicine against various objections.

However, the curative record of medicine, considered across all times and places is dismal. Yet medicine has persisted in many times and places. Why? I argue that this persistence shows that the business of medicine – the practice of a core medical competence – cannot be cure, even if that is the goal. Instead, what doctors provide is understanding and prediction, or at least engagement with the project of understanding health and disease. I argue that this Inquiry Model explains the persistence of ineffective medicine, and defend it against various objections.

6B. Pre-Empting Cures for the Inquiry Model of Medicine

Chad Harris

In this presentation I argue that Alex Broadbent's Inquiry Model of medicine (Broadbent 2019), which hypothesizes that medicine's core competence is inquiry even if its main goal is cure, improves our understanding of medicine but lacks a convincing account of the link that should exist between inquiry and cure. I take as the starting point of my argument two issues Broadbent raises in motivating for the Inquiry Model of medicine over the Curative Model, namely: 1) the persistence of medicine despite the lack of effective cures and 2) the continuing popularity of 'alternative' medical traditions in environments where so-called 'mainstream' medicine is readily available. I argue that Broadbent's reliance on the notion of inquiry to explain medicine's persistence in the face of 1) is undermined by considerations related to 2). I argue that the Inquiry Model's characterization of medicine's core competence as "engagement with the project of understanding health and disease" makes it difficult to explain why certain traditions become extinct and others persist. It also makes it difficult to explain the continued dominance of mainstream medicine over other traditions.

I suggest that the only way to judge whether a tradition satisfies its core competences and engages with the project of understanding health and disease, and by extension to explain why traditions die while others thrive, is by bringing in considerations related to cure. I further motivate for the division of medicine's core competencies into outward-looking, or patient-centric, and inward-looking, or profession-centric competencies. I endorse the Inquiry Model's conclusions related to understanding health and disease when it comes to inward-looking competencies, but I maintain that when it comes to patient-centric competencies the Inquiry Model falls short. I make the case that what patients seek from medical consultation is less about a deeper understanding of their condition and more about seeking expertise about ameliorating real or perceived vulnerability to threats they recognise as beyond their control. This explanation is consistent with the persistence of medicine despite the lack of effective cures, as well as with the persistence of alternative traditions alongside mainstream medicine. I conclude that inquiry and the pursuit of cure are both integral components of medicine, but the inquiry model's account of their relationship needs some adjusting.

Broadbent, A. (2019). *Philosophy of Medicine*. Oxford University Press.

6C. Medicine is Not Science

Olaf Dammann

The term "medicine" is variably used to mean medication, what doctors do, what medical students learn, or even the entire social endeavor of restoring and preserving health. Sometimes it is said that "medical research" has yielded some sort of breakthrough discovery. In this paper, I argue that if there is such a thing as "medical science" it is not the same as "medicine", simply because useful medical knowledge is not generated in medicine or by medicine, but for medicine. Medicine is a field whose practitioners apply a set of information management and communication rules with the goal of preserving or improving individual health. It is the practice of health interventions based on etiological and prognostic knowledge that is, in turn, generated by scientific research in biology, pharmacology, epidemiology, biostatistics, psychology, sociology and other scientific disciplines.

Broadbent (2019) thinks that “the goal of medicine is to heal the sick, but the core business of medicine is to understand and predict health and disease.” Although I agree with the first part of Broadbent’s inquiry thesis, which views medicine “fundamentally as an inquiry with a purpose”, I am not so sure about the second. While medical inquiry may be geared towards the generation of some kind of knowledge, that knowledge is not generalizable. Although medical practice may be capable of generating knowledge about individual patients and the effectiveness of interventions in these individuals, such knowledge is not generalizable, precisely because it is built on data from individuals, not populations, and thus not helpful to medical practice beyond the case at hand. What would medicine look like if doctors had only useless knowledge at their disposal? Even evidence-based medicine (EBM) is not science, but the application of knowledge from research synthesis and meta-analysis. EBM may perhaps count as meta-science, but has itself been criticized for excluding observational research and mechanisms, and for employing an evidence hierarchy that is not evidence-based.

At best, thinking of what medicine does as “medical science” is a harmless misnomer. At worst, it would suggest that the practice of medicine can yield generalizable knowledge, which would give case reports the power of randomized trials. This would be properly called “experience-based medicine”, which is an eighth alternative to EBM in addition to the seven proposed by Isaacs & Fitzgerald (1997). Indeed, the widespread use of such approach would catapult medicine back by decades, if not centuries.

Broadbent, A. (2019). *Philosophy of Medicine*. Oxford University Press.
Isaacs, D, Fitzgerald, D. (1999) Seven Alternatives to Evidence based Medicine. *British Medical Journal* 319.

6D. What is Modern Medicine?
Jonathan Fuller

Historians sometimes suggest that ‘modern medicine’ arrived by around the turn of the Twentieth Century, and in intellectual histories the transition to modern medicine is often described as the development of ‘scientific medicine’ (for instance, see Bliss 2011). What does it mean to say that modern scientific medicine began by this time? What is modern scientific medicine? What made the medicine of this period newly scientific? It was not the emergence of the medical sciences (anatomy, pathology, physiology, and so on) that made medicine newly scientific because these medical sciences had arrived earlier and had not at first had any great impact on medical practice. Rather, I argue that what made medicine modern and scientific was its adoption of the rationality of the medical sciences. Medicine began to explain and classify disease using concepts and theory from the medical sciences; it began to see the goal of medicine as intervening in these diseases; and it began to apply scientific (experimental, theoretical and above all empirical) reasoning in diagnosis and treatment. Syphilis in 1907 was an infectious disease caused by a specific germ (Spirochaete pallida); the doctor’s job was to cure the disease by eliminating the infection; and they attempted to do so using microbiological/physiological theory as well as trial and error reasoning (Osler & Churchman 1907).

However, this is not a fully satisfying account of modern scientific medicine because claiming that modern scientific medicine is medicine that adopts the rationality of medical science says nothing about what makes for medical science. What distinguishes early Twentieth Century medicine from Hippocratic medicine, which was similarly naturalistic and theoretical and sometimes even recommended intervening in disease? The problem is one of demarcation: what demarcates scientific medicine practiced in the early 1900s from pseudoscientific medicine (e.g. homeopathy or Hippocratic medicine)? In answering this question, I will not provide a general scientific demarcation criterion, but a more local and historical demarcation feature for scientific medicine: what distinguished scientific medicine from pre-scientific medicine was that the former but not the latter revised theory and practice in light of evidence, both seeking evidence and responding to it. Compared to earlier medical traditions, modern scientific medicine sought evidence (e.g. through dissection and laboratory study) and updated theory and practice in response (finally rejecting humourism and bloodletting in the Twentieth Century). This hallmark of science marked a turning point in the history of medicine, and could be used to demarcate scientific medicine from alternative medicines even today. Despite the rise of so-called 'evidence-based medicine' only around the turn of the Twenty-First Century, what made medicine modern 100 years earlier was its embracement of science and evidence.

Bliss, Michael. 2011. *The Making of Modern Medicine: Turning Points in the Treatment of Disease*. University of Toronto Press.

Osler, William and John W. Churchman. 1907. Chapter XII. Syphilis. In *Modern Medicine: Its Theory and Practice*, ed. W. Osler and T. McCrae, 436-521. Philadelphia and New York: Lea Brothers & Co.

6E. What is Personalised Medicine?

Zinhle Mncube

To personalise is to tailor or design something to meet an individual's needs within a specific context. Each time I watch a movie or rate a series on the platform, Netflix uses this data to personalise my viewing experience. That is, the particular view that I have of the content on Netflix is adapted to my current interests. What I am concerned with in this talk are debates in the biomedical literature about how to personalise medicine.
At its most basic level, personalised medicine (PM) means to tailor medical interventions to an individual patient's needs in a way that is effective for medical goals. As the popularity of PM has grown, there have been various attempts in the biomedical literature to define the term. But not only does 'personalised medicine' go by many names – genomic medicine, stratified medicine, precision medicine, amongst others. There exist various interpretations of the term.

For some, PM is targeted treatment to populations of patients stratified into subgroups according to their response to a particular drug as a result of a shared genetic variant. For others, the 'personal' and/or 'person' in PM refers to tailoring health care in a holistic manner according to an individual patient's needs (including, for example, her spiritual beliefs). Others hold what I call a broad-based conception of PM where several characteristics of an individual (biological, psychological, environmental, genetic and non-genetic disease predictors, for example) are used to determine and guide medical and healthcare decisions.

Yet others hold what I call a biological and pathophysiological conception of PM where primarily the biological characteristics of an individual patient (like genes) are used to determine and guide medical and healthcare decisions. These "some" or "others" are part of the 'PM debate' – they are stakeholders such as pharmaceutical and biotech companies, clinicians, doctors, scientists, academics, and patient advocacy groups, who debate the nature and scope of PM.

In this talk, I describe and evaluate biological and pathophysiological conceptions (or PBCs) of PM. I propose that there is a basic, core concept behind PM and that we should conceive of and assess the different conceptions of PM in the scientific literature as different forms or ways of achieving this personalisation of medicine. I argue that we should understand arguments against PBCs of PM as questions of the extent to which these conceptions truly personalise medicine to the individual patient. I contend that because PBCs are largely gene-centric, outside of single-gene disorders, PBCs are impoverished ways of personalising medicine to the many genetic and non-genetic factors that influence variability in health and common diseases among patients.

7. Epistemico-ethical Aspects of Animal-based Research

Sara Green and Simon Lohse

In recent years, animal experimentation practices have received much attention in philosophy of science, bioethics, sociology of science, and related fields of research. Examples include philosophers of science studying issues of representation and extrapolation in animal modelling (e.g. Baetu, 2014, 2016), bioethical assessments of chimera and xenotransplantation research (e.g. Hyun et al., 2007; Jorqui-Azofra & Romeo-Casabona, 2012), and sociological analyses of the public perception of laboratory animal research (e.g. von Roten, 2013). In these and similar cases, epistemic/methodological issues and ethical/social issues are typically analysed in isolation from each other. Philosophers of science, for instance, tend to focus on epistemic and methodological aspects of animal experimentation, while bioethicists usually restrict their analyses to moral and social issues. This can be problematic, as social and epistemic issues, and/or ethical and epistemic aspects, seem to be interconnected in many animal-based research practices.

This observation is reflected, to some degree, in recent work in philosophy of science. Ankeny and Leonelli's (2011, 2016) and Hardesty's (2018) assessments of model organism research in the life sciences, for instance, explicitly draw on perspectives from science & technology studies and philosophy of science to analyze the interconnectedness of epistemic and social functions of model organisms. It is noteworthy, however, that an explicit discussion of the (often intertwined) epistemic and ethical aspects of animal-based research is somewhat rare. This symposium aims at addressing this desideratum via an integrated discussion of animal-based research practices. Bringing together perspectives from philosophy of science (in practice), research on the ethical, legal and social issues of the life sciences ("ELSI research") and science & technology studies, the contributions to this symposium will explore aspects of animal experimentation practices where epistemic and ethical issues (broadly construed) are intertwined.

Simon Lohse's talk will sketch the landscape of use of experimental animals and non-animal methods, and discuss institutional as well as epistemic factors that currently provide barriers to development and use of non-animal-methods against the backdrop of regulatory/ethical challenges. Lohse will argue that the inertia in model replacement in biomedicine is rooted in the socio-epistemic logic of experimental science, as well as in secondary epistemic functions of animal models. Nicole Nelson's talk will discuss procedures and trade-offs in model refinement, specifically in attempts to increase the reproducibility of experimental results through blinding. Her examples illustrate how considerations about how to increase data quality are not straightforwardly related to blinding, but also to considerations about caring for experimental organisms. Sara Green's talk will document how both aims are pursued via patient-derived models in precision oncology. 3D organoids and patient-derived xenografts developed from human cancer tissue are promoted as models that can replace and refine standard murine models in cancer research and drug testing. The current hope is, however, also confronted with new epistemic as well as ethical challenges.

The talks will, hence, highlight the entanglement of socio-political contexts and animal-based research, and analyze the interplay of epistemic and ethical aspects in scientific practice and decision making.

Ankeny, R. A., & Leonelli, S. (2011). What's so Special About Model Organisms? *Studies in History and Philosophy of Science Part A*, 42(2).
Ankeny, R. A., & Leonelli, S. (2016). Repertoires: A Post-Kuhnian Perspective on Scientific Change and Collaborative Research. *Studies in History and Philosophy of Science Part A*, 60.
Baetu, T. M. (2014). Models and the Mosaic of Scientific Knowledge. the Case of Immunology. *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences*, 45.
Baetu, T. M. (2016). The 'Big Picture': The Problem of Extrapolation in Basic Research. *The British Journal for the Philosophy of Science*, 67(4).
Hardesty, R. A. (2018). Much ado about Mice: Standard-setting in Model Organism Research. *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences*, 68–69.
Hyun, I., Taylor, P., Testa, G., Dickens, B., Jung, K. W., McNab, A., … Zoloth, L. (2007). Ethical Standards for Human-to-Animal Chimera Experiments in Stem Cell Research. *Cell Stem Cell*, 1(2).
Jorqui-Azofra, M., & Romeo-Casabona, C. M. (2012). Some Ethical, Social, and Legal Considerations of Xenotransplantation. *Methods in Molecular Biology*, 885.
von Roten, F. C. (2013). Public Perceptions of Animal Experimentation across Europe. *Public Understanding of Science*, 22(6).

7A. Scientific Inertia in Replacing Animal Experimentation in Biomedicine
Simon Lohse

Millions of animals are used for basic and translational research purposes every year. In 2011 alone, more than seven million vertebrates and cephalopods were used in Europe. To reduce this number and to comply with the 3R-principle ("refine, reduce, replace"), animal rights activists, politicians and sympathetic life scientists have been promoting the development and use of alternative methods to animal experimentation based on cell and tissue cultures, organ(s)-on-a-chip technology and computer modelling (in short "non-animal-methods"). These efforts have, however, not led to an extensive replacement of animal experimentation in basic and translational science. In this talk, I will attempt to shed some light on this state of affairs with reference to key institutional and socio-epistemic barriers for the development and implementation of non-animal methods in the context of biomedicine.

In the first part of my talk, I will provide some background regarding animal experimentation in Europe and sketch the current landscape of non-animal-methods. I will then highlight a number of institutional factors that inhibit the development and use of non-animal-methods, such as the current funding structure and challenges in establishing the right kind of (data-) infrastructure. In the main part of my talk, I will turn to a socio-epistemic issue that has received some attention in the literature, namely the relatively low level of engagement of the scientific community in developing and promoting non-animal-methods. This situation is usually accounted for in two contrasting ways. The first way is based on the assumption (shared by many basic researchers) that animal experimentation is just indispensable for progress in biomedical science. Others (usually developers of non-animal-methods and activists) state that it is mainly dogmatism that inhibits the development and use of non-animal-methods. Both accounts, while containing some truth, fall short of explaining the complexity of the situation. For this reason, I will develop an alternative and more sophisticated explanation for the relatively low level of engagement of the scientific community which is based on insights from philosophy and sociology of science. More precisely, my talk draws on recent work on model organism research and scientific repertoires (Levy and Currie 2015; Ankeny and Leonelli 2016) and on the so-called "risk-spreading-argument" (Kuhn 1959/D'Agostino, 2010). I will argue that the inertia in replacing animal experimentation in biomedicine is partly rooted (a) in secondary epistemic functions of animal-based systems of practice (such as anchoring research communities and establishing shared methods and standards) and (b) in the socio-epistemic logic of science in general. I will show that my account offers a deeper explanation of the relatively low level of engagement of the scientific community that can integrate the true aspects of the discussed "standard accounts". In the last part, I will integrate my socio-epistemic analysis with ethical considerations in order to highlight the normative complexity of the regulation of animal research and inevitable epistemico-ethical trade-offs.

D'Agostino, Fred. 2010. *Naturalizing Epistemology: Thomas Kuhn and the "Essential Tension."* Basingstoke: Palgrave Macmillan.

European Commission. 2013. *Seventh Report on the Statistics on the Number of Animals used for Experimental and other Scientific Purposes in the Member States of the European Union SWD(2013)497*. Brussels.

Ankeny, Rachel A., and Leonelli, Sabina. 2016. Repertoires: A Post-Kuhnian Perspective on Scientific Change and Collaborative Research. *Studies in History and Philosophy of Science Part A*, 60.

Kuhn, Thomas S. 1959. The Essential Tension: Tradition and Innovation in Scientific Research. In: C. Taylor (ed.): *The Third University of Utah Research Conference on the Identification of Scientific Talent*. Salt Lake City: University of Utah Press, 162–174.

Levy, Arnon, and Currie, Adrian. 2015. Model Organisms are Not (Theoretical) Models. *The British Journal for the Philosophy of Science*, 66(2).

7B. Blinding as an Epistemico-ethical Decision

Nicole C. Nelson

The vast majority of researchers working with animal models do not report blinding their experiments. Systematic reviews of animal studies show that only 14-24% of publications report blinding the investigators or outcome assessors (Hackam and Redelmeier 2006; Kilkenny et al. 2009). Bias from lack of blinding is thought by some researchers to be an important driver of irreproducibility and translational research failures. Animal experiments that are not performed blind tend to report larger effect sizes, larger odds ratios, and smaller p-values than studies that are blinded (Bello et al. 2014; Holman et al. 2015). Under what conditions do researchers come to see blinding as a key component of good scientific practice, or not?

This talk will examine how "care for the data" (Fortun and Fortun 2005) and care for the animals are both implicated in decisions about blinding. Proponents of blinding make their case in both epistemological and ethical registers: they argue that blinding not only produces more robust, replicable data; it also prevents animal lives from being "wasted" in experiments that contribute little to the research record. Those who are more skeptical of the value of blinding likewise portray it as an epistemico-ethical decision:

they argue that blinded experimenters or care staff will be less attentive to their animals, compromising both the quality of data collection and of animal care.

These entanglements between epistemic and ethical concerns exist at the institutional level as well as the argumentative level. Despite the supposed separation of academic research and animal care into different committees and labor forces, changing research practices requires coordination across these domains. Researchers cannot, for example, unilaterally decide to assign a random code to indicate which treatment their mice have received if those mice are housed in a central animal care facility with its own practices and policies. Thus, as a matter of practical concern, those seeking to change research practice to enhance reproducibility must address ethical as well as epistemic considerations to succeed in shifting practices.

Bello, Segun, Lasse T. Krogsbøll, Jan Gruber, Zhizhuang J. Zhao, Doris Fischer, and Asbjørn Hróbjartsson. 2014. "Lack of Blinding of Outcome Assessors in Animal Model Experiments Implies Risk of Observer Bias." *Journal of Clinical Epidemiology* 67 (9).
Fortun, Kim, and Mike Fortun. 2005. Scientific Imaginaries and Ethical Plateaus in Contemporary U.S. Toxicology. *American Anthropologist* 107 (1).
Hackam, Daniel G., and Donald A. Redelmeier. 2006. Translation of Research Evidence From Animals to Humans. *JAMA* 296 (14).
Holman, Luke, Megan L. Head, Robert Lanfear, and Michael D. Jennions. 2015. Evidence of Experimental Bias in the Life Sciences: Why We Need Blind Data Recording. *PLOS Biology* 13 (7).
Kilkenny, Carol, Nick Parsons, Ed Kadyszewski, Michael F. W. Festing, Innes C. Cuthill, Derek Fry, Jane Hutton, and Douglas G. Altman. 2009. Survey of the Quality of Experimental Design, Statistical Analysis and Reporting of Research Using Animals. *PLOS ONE* 4 (11).

7C. Organoids and Mouse Avatars in Precision Oncology

Sara Green, Mie Seest Dam & Mette Nordahl Svendsen

A common translational problem in oncology is that cancer tumors growing in the body of individual patients often respond differently to medical treatments, compared to standardized cancer cell cultures or mouse models. The translational gap is often explained with reference to the limitations of standardized models to capture the genetic heterogeneity observed in the cancer clinic. To account for such limitations, patient-specific models are currently being developed from human tumor tissue. These include 3D cultures called "tumor organoids" as well as patient-derived xenografts (PDX), informally often called "mouse avatars". Patient-specific models are hoped to bring about three major breakthroughs in medicine and biomedical research.

First, they facilitate new epistemic and collaborative potentials (Davies 2012), e.g., for development of so-called living biobanks of cryopreserved, but viable, tumor cultures with specific mutational profiles. Such biobanks could be an invaluable experimental resource for cancer research to refine and supplement traditional experimentation on cancer cell lines or standardized mouse models. Second, it is hoped that personalized models break way for a "one patient paradigm" by providing better evidence for how individual patients will respond to targeted treatments (Akkerman and Defize 2017; Malaney et al. 2014). Third, patient-specific models are hoped to be able to replace some of the experimental animals currently used for drug development and drug testing (Jackson and Gareth 2017). The latter expectation is not only based on the potential for human tissue models to replace animal models, but also on the hope that development of better animal models (PDXs and humanized mouse models) will change the requirements of large sample sizes of test organisms in drug trials.

The hope and hype of patient-derived models are, however, often confronted with scientific uncertainties and ethical dilemmas. By combining insights from philosophy of science and ethnographic studies, we explore how such challenges materialize in concrete research projects and clinical practices involving PDXs and organoids. While patient-derived models are aimed to close the translational gap between human target and experimental models, they do not straightforwardly facilitate more certain inferences from model to target. In addition to questions about the extent to which patient-derived models represent robust models of specific patient tumors, important challenges include practical difficulties of growing human tissues outside a human body, as well as the concern that increasing the heterogeneity of cancer models decreases the potential for reproduction and extrapolation of findings. These epistemic challenges are highly intertwined with ethical issues in the cancer clinic, where scientific uncertainty must be balanced with considerations about care for current and future animals and patients. Hence, the epistemic implications of organoids and mouse avatars cannot be separated from their social and ethical functions.

Akkerman, Ninouk, and Libert HK Defize. 2017. Dawn of the Organoid Era: 3D Tissue and Organ Cultures Revolutionize the Study of Development, Disease, and Regeneration. *Bioessays* 39(4).
Davies, Gail. 2012. What is a Humanized Mouse? Remaking the Species and Spaces of Translational Medicine. *Body & Society* 18(3-4).
Jackson, Samuel J., and Gareth J. Thomas. 2017. Human Tissue Models in Cancer Research: Looking Beyond the Mouse. *Disease Models & Mechanisms* 10.
Malaney, Prerna, Santo V. Nicosia, and Vrushank Davé. 2014. One Mouse, One Patient Paradigm: New Avatars of Personalized Cancer Therapy. *Cancer Letters* 344(1).

8. Engineering and Biological Realizability

Mahi Hardalupas

Philosophers of science are increasingly interested in the role of engineering in emerging scientific disciplines. While typically engineering has been distinguished from basic science, many have problematized the dichotomy between basic science and engineering, suggesting ways for them to inform each other (Knuuttila & Loettgers, 2013; Boon, 2011). Investigating biological engineering can provide new insight into the philosophy of scientific practice by taking seriously how engineering practices can influence epistemological and metaphysical debates within science. Indeed, several philosophers have argued that these practices play a role in informing basic science through causal explanation (Baxter, 2019), natural kinds (Kendig & Bartley, 2019) and multiple realization (Koskinen, 2019). This session builds on this work and examines cases from engineering-heavy biological sciences such as synthetic biology and computational neuroscience to explore their philosophical implications linked to realizability.

One hallmark of engineering-based practices is that they allow scientists to produce models and systems that would not come into existence naturally. This emphasizes the importance of realizability and modal claims to scientific reasoning and raises several new questions in connection to engineering. How do engineered artifacts and how-possibly explanations factor into scientific reasoning? In what ways can multiple realization be useful to scientists in understanding biological function? Janella Baxter will assess whether these realized possibilities developed in synthetic biology can still be explanatory for more traditional biological disciplines. Rami Koskinen will explore how biological engineers can empirically

test realizers and what determines which count as scientifically interesting for multiple realization. Mahi Hardalupas will show how multiple realization manifests in computational neuroscience to analyze claims that engineered computational models like deep neural networks can be understood as artificial model organisms. Catherine Kendig explores how new technologies, in addition to providing opportunities for the reengineering of biological systems in synthetic biological practice, also usher in new articulations of synthetic parthood and kindhood that may have wider metaphysical and epistemic impacts.
By showcasing some of the exciting ways that engineering-based practices can contribute to the philosophy of science in practice, we hope to create a space for discussion of biological engineering and encourage further work on the philosophical implications of engineering-based scientific practices in other disciplines.

Boon, M. (2011). In Defense of Engineering Sciences: On the Epistemological Relations between Science and technology. *Techné: Research in Philosophy and Technology*, 15(1).
Baxter, J. (2019). How Biological Technology Should Inform the Causal Selection Debate. *Philosophy, Theory, and Practice in Biology*, 11(2).
Kendig, C., & Bartley, B. A. (2019). Synthetic Kinds: Kind-making in Synthetic Biology. In J. Bursten (ed.) *Perspectives on Classification in Synthetic Sciences: Unnatural Kinds.* London: Taylor & Francis (pp. 78-96).
Knuuttila, T., & Loettgers, A. (2013). Basic Science through Engineering? Synthetic Modeling and the idea of Biology-inspired Engineering. *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences,* 44(2).
Koskinen, R. (2019). Multiple Realizability as a Design Heuristic in Biological Engineering. *European Journal for Philosophy of Science*, 9(1).

8A. What's Synthetic Biology Got to do with It? How Engineered Biomolecules have Explanatory Value

Janella Baxter

The emerging discipline of synthetic biology is characterized by the entrance of scientists from other disciplines – notably electrical engineering, computer science and chemistry – into bioengineering. One thing that distinguishes some research programs in synthetic biology from other programs in traditional areas of the life sciences is that the former has created biological entities that have not evolved and are unlikely to evolve were it not for human engineering efforts. For philosophers and scientists who claim that the explanatory focus of biology is what has actually evolved by natural means in our world, synthetic biology presents a puzzle (Mitchell 2008, 2009, 2015; Ronai 2017; Wose 2004). As the historian and philosopher of science, Evelyn Fox Keller (Keller 2009), has asked: Do the products of synthetic biology fall within the explanatory scope of more traditional areas of biology or do they merely expand the universe of entities to explain? In this talk I argue that some of the technologies and techniques developed by synthetic biologists have played an ineliminable role in the explanations biologists have formulated about functional structure of proteins. By examining the experimental uses synthetic technologies have been used in protein modeling, I draw a distinction between assay technologies and interventionist technologies. Assay technologies serve to make observable an otherwise unobservable process; whereas, interventionist technologies don't. Instead, interventionist technologies serve to break causal relationships in complex systems. Because the observations assay technologies generate are often of explanatory significance to biologists, phenomena that has not actually evolved by natural means are often integrated into biological theories and hypotheses.

8B. Could We Really Be Made of Swiss Cheese? Or, An Engineering Epistemology for Biological Realization.
Rami Koskinen

According to Hilary Putnam's (1975) popular slogan for the functionalist program in the philosophy of cognition "we could be made of Swiss cheese and it wouldn't matter". This idea, which relied on an analogy to the many-to-one relationship between computer hardware and software, has become known as the multiple realizability thesis (MRT). What ultimately matters in understanding cognition and other biological functions is not the particular material building blocks, because in principle there should be many ways to build systems that exhibit such capacities.

However, in recent years MRT has received a considerable amount of criticism (e.g., Bechtel and Mundale 1999; Polger and Shapiro 2016). Polger and Shapiro have noted that many philosophers have been ready to commit to hypothetical cases of MRT without there being any direct proof of their feasibility. The criteria for biologically relevant realization has also been challenged. How can we know that, e.g., a hypothetical replacement of neural circuits with silicon-based ones would count as a biologically interesting change of physical realizers? If the mechanism of electrical signaling is similar in both cases, it is unclear whether we have a genuine case of multiple realization (Shapiro 2000).

The paper draws epistemological lessons from recent developments in xenobiology and chemical engineering to address the problematic status of biological realization. As our ability to manipulate the material basis of living systems – Putnam's "stuff we are made of" – begins to radically increase due to technological advancements, we are in a position to empirically test these philosophical ideas. Not only can biological engineers probe previously unrealized biological realizers, they can also control the intended functions of the systems they study (Koskinen 2019). The paper builds on previous work to specifically address the issue of what kinds of realizers count as "scientifically interesting" from the point of view of functionalist ideas.

8C. Engineering Artificial Model Organisms: A Role for Deep Neural Networks in Neuroscience?
Mahi Hardalupas

Recently, there has been growing enthusiasm surrounding the use of deep neural networks (DNNs) in neuroscience. However, since DNNs are only loosely inspired by biological neural networks and are mostly engineered for purposes other than understanding the brain, this raises questions of what their role is in neuroscience. Some researchers have suggested that deep neural networks are 'artificial model organisms' that populate a menagerie of 'DNimals' which serve a similar function to animal model organisms (Musall et al, 2019; Scholte, 2018). In this paper, I start by interpreting both DNNs and animal model organisms as attempts by scientists to engineer multiple realizations of cognitive processes. Then, drawing on Ankeny and Leonelli's account of model organisms, I explore in what ways the use of DNNs in neuroscience bears similarities and differences to animal model organism research. Through doing so, I evaluate the role that deep neural networks can play in neuroscience.

Ankeny & Leonelli (2012) highlight several characteristics of model organism research that contribute to their broad representational scope and targets. Drawing on their account, I argue that, while DNN research shares some of these features, they differ in important respects, for example, in how the models are engineered. Animal model organisms are engineered aiming at genetic standardisation ensuring minimal variability across experiments. In contrast, deep neural networks rely on engineered behavioural standardisation, which can provide possible hypotheses on how tasks constrain neural systems and behaviour (Kell & McDermott, 2019). I show this has consequences for the inferences that can be drawn from these models reflecting the different goals underlying their use, where animal models are often tied to medical research and DNNs to cognitive research. This allows us to distinguish between the features involved in creating a good model organism for cognitive science from those for clinical research.

I end by linking these ideas back to the larger question of why scientists engineer multiple realizations of phenomena such as cognitive processes. Elucidating the reasons why scientists attempt to engineer certain phenomena and the process by which they achieve this can help us better understand the contribution that engineering-based practices make to scientific research more generally and ultimately provides a fuller picture of scientific practice.

Ankeny, R. A., & Leonelli, S. (2012). What is So Special About Model Organisms? *Studies in History and Philosophy of Science Part A,* 42(2).

Kell, A. J., & McDermott, J. H. (2019). Deep Neural Network Models of Sensory Systems: Windows onto the role of Task Constraints. *Current Opinion in Neurobiology*, 55.

Musall, S., Urai, A. E., Sussillo, D., & Churchland, A. K. (2019). Harnessing Behavioral Diversity to Understand Neural Computations for Cognition. *Current Opinion in Neurobiology*, 58.

Scholte, H. S. (2018). Fantastic DNimals and where to find Them. *NeuroImage*, 180(September 2017).

8D. Between Technology and Biology: How Kind-making Works in Synthetic Biology

Catherine Kendig

What are synthetic kinds and how might they be conceived of as kinds? If we conceive of a synthetic kind as a form of life (or at least a life-like thing) whose construction is the result of human-assisted engineering, we effectively define a synthetic kind in virtue of its origin. For instance, we might suggest that an E. coli population harboring a synthetic DNA plasmid is a synthetic kind because it consists of a combination of the wild-type host organism and a synthetic plasmid. But is this the best way to conceive of it? As the E. coli owes its existence to both engineered and native wild-type sources, its causal origins are actually the result of both engineered and naturally occurring evolved parts and processes.

Conceived of within synthetic biology (and more broadly within bioengineering), synthetic kinds do not fit perfectly within either what are generally referred to as engineered or 'technological kinds', nor do they fit with what are generally understood to be natural kinds (cf. Schyfter, 2012). If they are neither fully technological nor natural biological kinds, does this pose a problem for conceiving of these entities as kinds at all? Is there a third option? One alternative we might consider is that any determination of what kind of kinds exist in synthetic biology may be the sort of thing that ultimately rests on understanding the nature of the constructive processes employed within synthetic biology itself. In exploring this suggestion further, I aim to shed light on both the synthetic processes that are entailed in the making of material objects, mechanisms, processes, and pathways as well as the consideration of multiple causal origins that may be used in the determination of synthetic kindhood. These processes include, but are not limited to, the theoretical construction of models and algorithms and the devising of repeatable methods and techniques to delimit the subject of investigation, engineering, and other scientific and operational wanglings. Building on previous work (Kendig 2016), I suggest that these are the epistemic and ontological kind-making ('kinding') activities by which the categories of that discipline or subdiscipline are configured (e.g. through theoretical and computer-assisted modelling, engineering design, exploratory 'proof-of-concept' investigations, development of platform technologies, and consideration of the ethical, legal and social impacts of the generation of these categories). By investigating recent developments in agriculture and food science, I illustrate how synthetic biological research and the non-native products and processes that arise from it seem to suggest a notion of kindhood that does not rely on the pre-carved up contents of the world but instead relies on the ineliminable role of the carvers (e.g., the synthetic biology practitioners as kind-makers or kind-locators), as well as the ascribers and cataloguers of the joints partitioning the wholes.

Kendig, C. (2016). What is *Proof of Concept* Research and how does it Generate Epistemic and Ethical Categories for Future Scientific Practice? *Science & Engineering Ethics* 22(3).
Schyfter, P. (2012). Technological Biology? Things and Kinds in Synthetic Biology. *Biology & Philosophy* 27.

9. Research Projects as Contexts for Experimental Protocols, Epistemic Activity, and Modelling in Biology

Robert Meunier and Steve Elliott

The category of a research project lends itself well to the analysis of scientific practice (Meunier 2019). It can be defined as a historical episode in which a research team addresses questions or problems about phenomena according to a methodology, and by using a set of materials, devices, and modelling and representational techniques. Projects instantiate experimental systems (Rheinberger 1997) or observational regimes, and they are often designed with, and later contribute to, the materials, theories, norms, and practices assembled in repertoires, which projects can on occasion seed (Leonelli and Ankeny 2016). Importantly, the constituents of a project are assembled according to a plan designed from the original question (Meunier 2019). Researchers revise projects and their plans dynamically in response to unexpected findings or failures, and they develop and redeploy specialized knowledge about how to conduct projects. Ultimately, the category of project helps explain how the components of experimental systems are transferred and how repertoires are instantiated, altered, and seeded.

Philosophers have only recently focused on projects, and further development of the category requires detailed studies of how practices and projects mutually influence each other. We propose a session focused on the theme of how material practices of experimentation and theoretical practices of modelling can be better understood from the perspective of the projects that engender them. These practices can then be described and explained partly as taking shape in the context of dynamically adjusted goals, methodologies, and evolving research landscapes.

The papers discuss case studies from developmental biology and synthetic biology, importantly contributing to the study of the epistemic structures and historical dynamics of those highly integrative fields (Love 2015; Knuuttila and Loettgers 2014). The papers develop analytic tools to analyze how projects and practices influence each other in those, and potentially other, fields.

Steve Elliott shows how a project’s aims influence researchers as they combine, retool, and iterate standardized research protocols so that regimented practices yield meaningful data. Equally interested in experimental practice, Robert Meunier shows how epistemic activities are directed at research objects to single out relations between them in a way that is structured by and instantiates a project’s goal. Analyzing a case from synthetic biology, Andrea Loettgers and Tarja Knuuttila show how new research projects are often prompted by the unexpected results generated by the practices of material and mathematical modelling involved in earlier projects pursued in a lab.

Ankeny, R. A., & Leonelli, S. (2016). Repertoires: A Post-Kuhnian Perspective on Scientific Change and Collaborative Research. *Studies in History and Philosophy of Science Part A,* 60.
Knuuttila, T., & Loettgers, A. (2014). Varieties of Noise: Analogical Reasoning in Synthetic Biology. *Studies in History and Philosophy of Science Part A,* 48.
Love, A. (2015). Developmental Biology. In E. N. Zalta (Ed.), *The Stanford Encyclopedia of Philosophy* (Fall 2015).
Meunier, R. (2019). Project Knowledge and its Resituation in the Design of Research Projects: Seymour Benzer’s Behavioral Genetics, 1965-1974. *Studies in History and Philosophy of Science Part A*, 77.
Rheinberger, H.-J. (1997). *Toward a History of Epistemic Things*. Stanford University Press.

9A. Protocols Tailored for Research Projects and for Integration of Practices
Steve Elliott

Philosophers increasingly study practices to help characterize the epistemology of science (e.g. Soler at el. 2014). They acknowledge that practices are often shaped by local research contexts, but to date they have largely focused on how labs shape practices, especially for training junior researchers and for solving problems (e.g. MacLeod and Nersessian 2016). These authors acknowledge that problem-solving and other aims can be fruitfully conceptualized as aspects of research projects conducted within labs, but doing so remains itself an open project. This task is hindered by the imprecise and amorphous concept of practice. How might we conceptualize practices so we can study how they are modified in relation to local project aims? This paper builds on previous work to address that question. I study practices as operationalized in research protocol documents, and I show how those documents change over the temporal span of a project.

I focus on the lab of biologist Greg Wray at Duke University, and on a project he conducted with his graduate students between 2010 and 2017 (e.g. Garfield et al. 2012). The team studied the development and evolution of purple sea urchins. The team modified two widely-used and standard protocols. The first is a breeding design protocol that yields offspring controlled for genetic and phenotypic variation. The second is an assay protocol to characterize the amount of RNA within a sample drawn from subsets of the population. Without modifying these protocols, the practices of Wray's team would not have yielded the data needed to address the research questions and epistemic aims that the team had set for their project. I show the standard protocols, how Wray's team modified those protocols, and how they used their resultant data to address their questions and aims. I suggest that this practice of protocol-tailoring is a common aspect of project research, and that it might be used to help distinguish projects from other categories of local contexts. Furthermore, as Wray's project studied the phenotypic effects of variation in sea urchins' gene regulatory networks, I suggest that protocol-tailoring enables a kind of integration of practices that can complement theory integration in evolutionary developmental biology.

Garfield, David A., Daniel E. Runcie, Courtney C. Babbitt, Ralph Haygood, William J. Nielsen, and Gregory A. Wray. 2013a. The Impact of Gene Expression Variation on the Robustness and Evolvability of a Developmental Gene Regulatory Network. *PLOS Biology* 11.
MacLeod, Miles, and Nancy J. Nersessian. 2016. Interdisciplinary Problem-Solving: Emerging Modes in Integrative Systems Biology. *European Journal for Philosophy of Science* 6.
Soler, Lena, Sjoerd Zwart, Michael Lynch, and Vincent Israel-Jost, eds. 2014. *Science after the Practice Turn in the Philosophy, History, and Social Studies of Science*. New York: Routledge.

9B. Epistemic Activities and Exemplification in Research Projects: Staining and Transplanting in Research on the Anatomy and Development of Nerves
Robert Meunier

The paper argues that the ways epistemic activities contribute to the production of knowledge (Chang 2014) can be understood better when seen in the context of research projects. Research projects are characterized and shaped by goals. More precisely, researchers typically draw on general knowledge about what can be known about things or on more specific knowledge about what can be known about the entities in the domain of phenomena they are concerned with. Such goal schemata are typically associated with knowledge about ways to achieve the goal (i.e. method schemata), again general or domain specific. Epistemic activities can now be analyzed in terms of the goal and method schemata that guide them.

Goal schemata might, for instance, specify what it means to know the composition of an entity (e.g. knowing its parts, their relative position, size etc.) or the causal relations in a domain etc. These goals suggest method schemata of decomposition or controlled intervention, respectively. Projects can be conceptualized as instantiating such method schemata, thereby making the material at hand exhibit the parts, properties or relations specified in the goal schemata. The key argument of the paper is thus that research activities aim to make the materials exemplify what satisfies the goal schema. Microanatomy and embryology between 1880 and 1920 are distinct but interrelated fields, characterized by different goal schemata, one concerned with the composition of organs the other with this composition as the outcome of a process of development. While microanatomy relies on practices of manual separation and staining of tissues as well as microscopy and camera lucida drawings, embryology needs to incorporate these techniques and combine them with strategies which enable researchers to track processes (Griesemer 2007) and analyze causal relations, such as transplantation experiments. The talk sketches two research projects, one from each field, to show how they are shaped differently by the goals and attendant methods and hence make the material exemplify different kinds of parts and relations.

Chang, H. (2014). Epistemic Activities and Systems of Practice. In L. Soler, S. Zwart, M. Lynch, & V. Israel-Jost (Eds.), *Science after the Practice Turn in the Philosophy, History, and Social Studies of Science* (pp. 123–150). Routledge.
Griesemer, J. (2007). Tracking Organic Processes: Representations and Research Styles in Classical Embryology and Genetics. In M. D. Laubichler & J. Maienschein (Eds.), *From Embryology to Evo-Devo: A History of Developmental Evolution* (pp. 375–433). MIT Press.

9C. Synthetic Biology and the Emergence of Novel Experimental Practices and Research Projects

Andrea Loettgers and Tarja Knuuttila

Analyzing a case from synthetic biology, Andrea Loettgers shows how new research projects are often prompted by the unexpected results generated by the practices involved in earlier projects pursued in a lab.

10. Homology and Development from Richard Owen to Developmental Genetics

Aaron Novick, Devin Gouvêa and Günter Wagner

'Homology' has been a central and controversial biological concept since Richard Owen first defined it in 1843 (and arguably even earlier, under other names), and it remains subject to substantial controversy even today. This symposium brings together work on the history, philosophy, and contemporary biology underlying these controversies, with an especial focus on interactions between the 'homology' concept and the study of development.

The first paper in the session, by Aaron Novick, focuses on Richard Owen's development of the 'homology' concept in the mid-19th century. The paper explores Owen's concept as an attempt to manage a diverse (and shifting) array of theoretical demands. 'Homology', for Owen, had to serve taxonomic ends (with the homology/analogy distinction serving as the successor to the much-debated affinity/analogy distinction), it had to conform to Owen's particular theory of animal development, and it had to interact with Owen's own evolutionary theorizing. The paper examines the strategies Owen used to shape the concept into a form flexible enough to serve these various ends.

The second paper in the session, by Devin Gouvêa, analyzes interactions among theory, concepts, and data in active research problems involving the homologies of vertebrate forelimbs. The paper focuses on how scientists have reconstructed the evolutionary transition from the five-fingered dinosaur hand to the three-fingered bird wing. Gouvêa shows that research on this transition has led to substantial change in both the internal structure of developmental theories of homology and in their relation to other theories. These changes have occurred under pressure from the expansion of molecular techniques and from the need to reconcile conflicts between competing perspectives.

Finally, Günter Wagner will present an account of homology in terms of traceability, using this to resolve long-standing questions about the sense in which homologous parts are "the same", as well as to connect the 'homology' concept to other important forms of traceability in biology (e.g., tracing molecules through the body). As part of this account, Wagner further develops a notion of homology of process, arguing that it is more fundamental than homology of structures: structures are homologous in a secondary sense, in virtue of the homology of relevant underlying processes.

In concert, the three papers contribute to a holistic appreciation of homology—both the concept and the phenomenon itself. They illuminate important aspects of the history of the concept, use the concept to explore issues in the philosophy of the conceptual structure and change, and help point the way toward the role of the concept in future biology. They do all of this while considering the various scientific practices that have shaped the concept over time: practices of taxonomizing, of reconstructing historical transitions, and of tracing.

11A. Reference and Ignorance: The Case of Richard Owen

Aaron Novick

This paper considers the question of how reference is fixed under conditions of ignorance, focusing on Richard Owen's attempt to delineate a workable 'ho-mology' concept. Owen developed the concept at a time when the nature of homology was unsettled and subject to vociferous dispute, and at a time when none of the available positions were, by modern lights, correct. It is in this sense that Owen was tasked with fixing the reference of 'homology' under conditions of ignorance.

Traditional philosophical accounts of reference fixation fall into one of two traditions: descriptive accounts and causal accounts. Descriptive theories of reference take reference to be fixed by description: a term refers to whatever entities or relations meet some appropriate description. Under conditions of ignorance, such theories predict that terms simply fail to refer, as nothing in the world will meet the descriptions given. I contend that Owen did successfully develop a workable, applicable notion of 'homology'. Descriptive theories of reference do not account well for this success. Causal theories of reference err in the opposite direction: they make successful reference under conditions of ignorance too easy. Causal theories of reference allow for reference to be fixed by an initial baptism, in which the baptizer stands in the appropriate sort of causal contact with the entity/kind/relation to which they refer. I contend, however, that the bit of the world that Owen was in "contact" with is too complex for reference to fix on any particular relation, as causal theories suggest that it should.

In this paper, I investigate the various strategies Owen employed to secure the reference of the term and to enable its successful use by a community of re-searchers who disagreed over numerous fundamental theoretical issues. They disagreed, for instance, over how to interpret the natural system, over the nature of the relations the natural system was supposed to capture, over the reality and nature of transmutation, over the relative priority of form and function, and more. 'Homology' played a role in all of these debates, yet Owen was able to successfully make the concept a shared resources that all parties could use to communicate effectively.

I show that Owen achieved this by balancing a diverse set of conceptual pressures. He developed an open-ended, flexible definition of 'homology' that could be preserved under various different background theories (and that in-deed still surfaces in contemporary discussions of homology as basically accurate). Alongside this, he developed criteria for homologizing parts of organ-isms that were generally acceptable, again across theoretical divides. Thus, the term could be successfully applied despite disagreement over the nature of the relation to which it referred.

I suggest that the best way to understand Owen's accomplishment here re-quires recognizing that, while the term could be applied definitely to particular cases (homologous parts could be clearly identified), the reference of 'homology' itself was importantly indeterminate—the term gestured, in a somewhat in-definite way, at a complicated, messy region of the world (a region whose complications largely remained to be discovered), and this very indefiniteness is crucial to Owen's success, as well as illuminating of the subsequent convoluted history of the term.

11B. Digit Homology and the Quest for Integration

Devin Gouvêa

During the evolution of modern birds, a five-fingered hand gave rise to a three-fingered wing. Comparative anatomical investigation of this transition have raised an enduring puzzle. To which of the five standard tetrapod fingers are the three avian fingers homology? Palentological and embryological investigations have consistently yielded different answers to this question. I identify two patterns that emerge from the recent history of scientific work on this problem. First, the conflict at its heart has been consistently characterized by involved scientists as a clash between different types of data. Attempts to resolve that conflict turn on the question of how such data should be interpreted to provide evidence for particular evolutionary scenarios. Second, progress on this problem has resulted from the repeated and reciprocal interplay between data gathering and hypothesis generation. Over time, new fossil finds and new developmental studies have repeatedly prompted the modification of existing evolutionary scenarios and the proposal of new ones. New scenarios have in turn encouraged the search for experimental data to confirm or challenge them.

These patterns both complement and complicate existing philosophical analyses of the concept of homology. Such analyses have primarily reckoned with the tensions between different theories of homological sameness: some scientists emphasize phylogenetic relationships, while others appeal primarily to developmental mechanisms. In the digit homology case, scientists specializing in paleontological and embryological data tend to fall on either side of this divide. However, their differences in theoretical perspective have not in themselves been a major source of tension or obstacle to scientific progress. Researchers recognize these differences but do not treat them as a problem to be solved or even as particularly critical to the resolution of digit homology. In response to this situation, I argue that the existence of competing (and even potentially incompatible) theoretical perspectives can be fruitful for the study of complex scientific phenomena. They provide important structure for the work of gathering, interpreting, and re-interpreting data that is necessary for integration to proceed.

11C. Traceability and Individuation

Günter Wagner

Traditionally, during the 20th century, the homology concept had a hard time finding its place among the foundations of modern biology. Here I will argue that it has to do with the conceptual structure of homology and suggest that recent developments in the philosophy of biological individuality and process ontology can lead to a clearer understanding of the nature of homology. These concepts are “traceability” and “individuation.”

Traceability is the ability to identify an object or an individuum across different observations of the same object or different instantiations of the same individuum. As such, tracing is a fundamental part of scientific practice but has not received formal consideration within the philosophy of biology or science in general. The notion of traceability is mostly developed in applied fields like software development, health care delivery, food technology and similar. Here I am taking this concept from James Griesemer (2019).Traceability in the sciences often includes a matching operation between a set of observations and a model of transmission or motion. For instance tracing a disease gene takes a family tree and the distribution of a heritable disease as observations and uses the Mendelian rules of segregation to reconstruct the most likely path of transmission among relatives. Similarly, homology statements are hypotheses about the traceability of a certain character across the phylogeny. The concept of traceability eliminates the necessity to link homology to a particular mode of inheritance or transmission. This is important as it has been argued that anatomical structures and cell types can not be homologous because there is no direct copying mechanisms involved, like it is the case of genes and DNA. The conceptual “essence” of homology is that it identifies features, body parts etc. that have a traceable history across the tree of life. Any entity that has such a traceable history is a candidate homologue, which in principle can be extended to functional processes as well as behavioral and even cultural traits.

For traceability to be apply to a feature, the candidate homologue has to be individuated from the rest of the body, genome, or among the cells of an organism. The need to assume individuation of a structure is an often overlooked feature of the homology concept. This feature is also important as it links the abstract notion of homology with mechanisms that enable a structure, gene or cell type to have its own quasi-independent evolutionary history.

As the notion of traceability is extendable to functional processes, this approach suggests a unitary grounding for complex anatomical and cell biological homologues. One may argue that the homology of these complex structures can be grounded in and tied to the homology of their individuating molecular processes. What exactly these individuating processes are is subject to ongoing research. Candidates are individuating gene regulatory networks like those suggested by Cobert for cell type identity, Kernels as suggested by Eric Davidson for developmental fields, and Character Identity Networks and Core Regulatory Complexes of transcription factors.

12. Making Geologic Time

Joeri Witteveen

The study of deep time is at the center of many sciences, including geology, evolutionary biology, and climate science. While exposed layers of rock around the world can provide a peek into the earth’s past, they far from deliver a readily legible record of our planet’s history. Geologists need to wrestle information about the chronology of the earth from the partial, perturbed, and incompletely preserved rock strata and their contents. In the second half of the 20th century, stratigraphers worldwide initiated an effort to construct a global Geologic Time Scale based on data about the lithologic, magnetic, chemical, biological, and other attributes of rock layers. This ongoing endeavor to construct a time-calibrated periodization of earth history provides fertile ground for a practice- based philosophy of geologic time. For example, the complex array of methods, models, and varieties of data involved in the measurement of geologic time provides new insights into the epistemic dynamics of calibration and correlation in the context of the historical sciences. Closely related to these themes about measuring geologic time are issues about the standardization and periodization of geologic time. The effort to construct a Geologic Time Scale prompts questions about the presuppositions and implications of the aim to formalize and unify the hierarchical units of geologic history. Some of these questions concern the epistemic and normative dimensions of conventionality and naturalness that are familiar from the philosophy of metrology and biological taxonomy, but that are manifested in the geologic context in new and unexpected ways. What is more, the governance of geochronology raises topical social epistemic questions about the relation between geology, other sciences, and society. This is illustrated in particular by the heated debate over the recognition of the ‘Anthropocene’ as geochronologic unit, which is thrusting questions about measuring and making geologic time into the limelight of the public eye.

12A. What is at Stake in the Formalization of a Chronostratigraphic Unit? A Case Study on the Anthropocene
Hernan Bobadilla

The Anthropocene Working Group (AWG) is preparing a proposal to formalize the status of the 'Anthropocene' as a new series/epoch. In order to be accepted, the proposal must meet the guidelines established by the International Commission on Stratigraphy (ICS). Rejection of the proposal is a realistic scenario, based on expert opinions pointing at procedural shortcomings of the proposal.

This paper examines what is at stake in the formalization of a chronostratigraphic unit, such as the 'Anthropocene'. More specifically, I assess the scientific (geological and non-geological) and non-scientific impacts of introducing the 'Anthropocene' as a formal unit. In order to conduct this assessment, I follow the next steps. First, I present the ICS's guidelines and discuss their motivations. Second, I present the core tenets of the AWG's proposal and discuss those aspects in which the proposal – in its current state – departs from the ICS's guidelines (including foreseeable recommendations). Third, I discuss the scientific and non-scientific motivations, as presented by the AWG, in introducing the 'Anthropocene' as a formal unit. Fourth, I consider the costs – both scientific and non-scientific – of the AWG's proposal being rejected.

In developing my argument, I construe the 'Anthropocene' case as one of tension between a unificationist agenda and a diverse landscape of scientific practices in the making. On the one hand, the unificationist side of this tension is embodied in the ICS. Indeed, one of the stated objectives of the ICS is to keep a unified nomenclature – or classification scheme – that serves as a global standard. The benefits of keeping a unified nomenclature are well-documented. Particularly salient are ease and effectiveness in communication, and institutional legitimacy. I submit that the AWG – by submitting their proposal to the ICS – validates the ICS's unificationist agenda and acknowledges the advantages of formalization. On the other hand, as a matter of fact, formalization has not been needed for scientists to employ the term 'Anthropocene'. To be sure, the term has been employed with dissimilar senses across various scientific disciplines, including the social sciences. That is, within each discipline (or research program) scientists establish their own standards. This state of affairs illustrates the diversity of scientific practices in the making.

In my assessment, I highlight three critical costs derived from the putative rejection of the AWG's proposal: i) ineffective communication among specialists from different scientific disciplines; ii) limited institutional legitimacy of scientific endeavours focused on 'Anthropocene' studies; and iii) failure to raise awareness of the impact of anthropogenic activity upon the Earth's dynamics. The first two problems have a direct impact in the productivity of scientific research. The third problem is characterized as a socio-political, with indirect and long-term impact in scientific productivity. The first problem is characterized as a problem of misunderstanding among scientists. The two latter problems have roots in an issue of public misunderstanding of science, namely in the idea that consensus and unified standards are landmarks of reliable science.

12B. Learning to Measure What Isn't There: The Problem of Missing Time
Alisa Bokulich

The primary source of our knowledge about geologic time and the Earth's 4.5 billion-year history is the stratigraphic (rock) record and the various different clues that are preserved in its layers. Although there are many interesting challenges in reconstructing geologic time from these records, one of the most difficult—and seemingly intractable—issues in the foundations of geologic time is what we might call the "problem of missing time." This problem results from gaps or "hiatuses" in the geologic record, which can arise either from stasis (no sediment deposited) or because the sedimentary layers once deposited were subsequently eroded away. Gaps appear in the stratigraphic record as an "unconformity"—a boundary between two different bodies of rock representing two discontinuous periods of time. The most famous of these is the Great Unconformity, first identified in the Grand Canyon (Powell 1875) and believed to represent anywhere from 100 million to 1 billion years of missing time. The Great Unconformity lies just below the Cambrian strata and its erosion history has been speculated to be a cause or effect of some of the most puzzling and important events in Earth's history (e.g., Snowball Earth, onset of plate tectonics, the rise of free O2 in our atmosphere, and the Cambrian explosion). How much time is missing from the geologic record? And precisely which periods of Earth's history do they represent? Answering these questions is essential not just for reconstructing geologic time, but also for beginning to discriminate among the above causal hypotheses. While one might have thought that these were intractable questions, whose answers were lost to time, surprisingly geoscientists are developing a new suite of methods, known as "deep-time thermochronology," to quantitatively measure the timing and duration of the rock record that isn't there.

The rise of deep-time thermochronology provides a striking example of what I call "unconceived opportunities" in the historical sciences, that is, the discovery of new sources of data about phenomena that we would have antecedently thought were not empirically accessible. These new sources of data are typically not "ready made" in the historical sciences, but rather require vast amounts of foundational laboratory work, field studies, and advances in modeling and theory to come together in order to extract this data, and turn "detritus into evidence" (Jeffares 2010). In this talk I analyze how geoscientists are learning to quantitatively measure the duration and timing of the gaps in the stratigraphic record using deep-time thermochronology. I draw three philosophical lessons from this case: First, there is far more experimental laboratory work that goes into the historical sciences than is often appreciated (e.g., Cleland 2001, 2002). Second, thermochronology provides a philosophically rich example of scientific measurement, advancing our understanding of so-called derived measurements (e.g., Parker 2017), model-based measurement (e.g., Tal 2012), and model-data symbiosis (e.g., Edwards 2010; Bokulich forthcoming). Third, and finally, I draw out the implications of this case for the "optimism vs. pessimism" debate about the historical sciences (e.g., Turner 2007, 2016; Jeffares 2010; Currie 2018).

12C. Golden Spikes, Silver Bullets, and the Ma(r)king of Chronostratigraphic Boundaries
Joeri Witteveen

The geologic time scale divides the history of the earth into a hierarchy of geochronologic (or chronostratigraphic) units, from eons down to eras, periods, epochs, and ages. Since the 1970s, the International Commission on Stratigraphy and its subcommissions have taken up the task of formally establishing the boundaries of these units through the designation of so-called “Global Stratotype Section and Points” (GSSPs). A GSSP sets a precise boundary between two geochronologic units and is marked by driving a “golden spike” into a rock section at a designated point.

A prima facie puzzling aspect of the GSSP approach is the recommendation that a golden spike be placed at a horizon where, geologically speaking, nothing happened. Even if a time unit was originally introduced on the basis of a perceived “natural break” in the stratigraphic record, the GSSP approach mandates that the unit's formal boundaries are to be placed in sections that lack any abrupt changes in lithology or fossil content. The rationale for assigning GSSPs to horizons of geologic non-events is to establish boundaries that do not shift under advances in stratigraphic resolution or changing perspectives on naturalness. GSSPs, it is often said, are theory-free.

Nevertheless, the GSSP approach continues to be criticized for presenting a mythical silver-bullet solution to the division of geologic time. Critics have taken issue with the GSSP approach on various grounds, and have variously claimed that it is methodologically inept, conceptually flawed, and even scientifically dishonest. In this paper, I will provide a philosophical take on the GSSP approach and its discontents. I argue that while there are patent opportunities to improve the process of designating and correlating GSSPs, the disapprobation of GSSPs as such is unfounded and tends to depend on a misunderstanding of GSSPs qua reference standards. In their efforts to describe and analyze the nature of GSSPs, stratigraphers have often relied on comparisons with biological type specimens or token-based measurement conventions. I argue that while GSSPs resemble both kinds of material standards in certain ways, a proper understanding of their functioning should focus on notable differences that bring their true nature into focus.

V. Contributed Talks

Abstracts of contributed talks are listed alphabetically by author surname.

1. Confirmation and Variety of Evidence in Multimessenger Astronomy

Shannon Sylvie Abelson

On August 17th, 2017 the LIGO and Virgo Observatories registered a gravitational wave signal which was then matched to a virtually simultaneous electromagnetic signal detected by the Fermi Gamma-ray Space Telescope. This ushered in the era of multimessengers, wherein astronomers search for coincident observations of the four main sources of galactic signal: gravitational rays, cosmic rays, neutrinos, and electromagnetic. The confirmation of models of these cosmic messengers depends heavily upon the variety of evidence thesis (VET). In short, where multiple and heterogenous sources of evidence epistemically converge upon the acceptance of a hypothesis, that hypothesis is made stronger than if its sources of evidence were homogenous. Recent philosophical treatments of the VET have questioned its formal legitimacy and effectiveness (Stegenga 2015; Claveau and Grenier 2019). And yet, many outstanding questions in astronomy require a multimessenger approach. In what follows I discuss the role variety of evidence plays in one of the most exciting discoveries connected to the 2017 detection event. I show that the convergence of multiple empirical observations of distinct phenomena emanating from the event are taken by the astronomical community to lend substantial confirmatory support to the thesis that the event was a kilonova.

Kilonovae are the emission scenarios resulting from the merger of compact objects, where the luminosity exceeds 1000 times that of a nova event. The luminosity generated is directly caused by the radioactive decay of light (e.g., silver) and heavy elements (e.g., gold) in a process called rapid neutron-capture nucleosynthesis, or the r-process. Theoretical kilonova models posit first that the event ejects high-velocity particles. The decay of nuclei leads to the production of light elements and causes a short-lived blue hue. This then develops into a red hue, indicative of the thermal process necessary for the formation of heavy elements. The origin of these heavy elements has been a significant missing piece in our understanding of the universe.

In a period spanning two weeks, each of the luminous phenomena associated with kilonovae were observed:

a) Gravitational Wave: the first detected signal, indicative of neutron star presence.
b) Optical: The luminosity of the event peaked with a blue hue, and then rapidly faded over several days.
c) Infrared: The infrared emission peaked in a range of 1-3 micrometers and lasted for two weeks.
d) Ultraviolet: the ultraviolet emission was short-lived and high-velocity.
e) Blackbody: the entities in the event were quasi-blackbody, indicating a thermal source (Metzger 2017).

Each of these emissions were detected at different wavelengths, using various instruments, and bear on different direct and indirect consequences of the kilonova model. Astronomers on the research teams in the kilonova case consistently and repeatedly make reference to the varied body of evidence. It is precisely the variety of the evidence that is accorded epistemic weight; the VET is clearly a methodological commitment in multimessenger astronomy. The dawn of multimessenger astronomy realizes a long-awaited goal of astronomical research. As such, the methodological commitments and practices are of particular significance to philosophical investigations of scientific work.

Claveau, François and Grenier, Olivier 2019. The Variety-of-Evidence Thesis: a Bayesian Exploration of its Surprising Failures. *Synthese* 198(8).
Metzger, Brian. 2017 Welcome to the Multi-Messenger Era! Lessons from a Neutron Star Merger and the Landscape Ahead. arXiv:1710.05931.
Stegenga, John 2015. Seeing Things: The Philosophy of Reliable Observation, review of Robert Hudson, *Seeing Things: The Philosophy of Reliable Observation*, Oxford University Press, 2014. *Notre Dame Philosophical Review*.

2. From Neuroimaging to the Clinic – Weighing the Uncertainties in New Diagnostics of Consciousness

Lise Marie Andersen, Hanne Bess Boelsbjerg and Mette Terp Hoejbye

Uncertainty is a basic condition of medicine - in research, in the clinic and in the intersection between the two domains. New technologies are continuously developed in order to make diagnostics more precise, however, with technological innovations uncertainties are often enhanced or transformed rather than eliminated. Researchers and physicians are thus presented with new challenges weighing the transformed uncertainties of the new context.

Our study traces the uncertainties involved in the development of neuroscientific tools for diagnostics in unresponsive patients with uncertain consciousness due to anoxic brain injury after an out-of-hospital cardiac arrest (OHCA). Most people who have survived an OHCA are initially comatose. While some may gain capacity of functional communication to different degrees, others never recover. The prognosis for recovery is radically different depending on whether the patient develops to a vegetative state (VS) or a minimally conscious state (MCS). Patients in VS are regarded as lacking the capacity for consciousness entirely, while patients in MCS are considered to have some minimal capacity for consciousness. Currently the standard method for determining level of consciousness is the JFK Coma Recovery Scale–Revised. The scale relies on behavioral cues, such as reaction to stimuli. Researchers have argued that this method is seriously flawed and that the rate of misdiagnosis and the related problem of prognosis is unacceptable. This has led to a wish for new tools of diagnostics and thereby more accurate prognostics for this subgroup of patients. Working to break new ground in diagnostics for these patients, a current experimental protocol in an interdisciplinary Danish research group is seeking to advance the functional magnetic resonance imaging (fMRI) method, in combination with electroencephalography (EEG) and other measures, as aide for making prognosis and the necessary medical decisions. Applying an interdisciplinary approach combining anthropological methods of semi-structured interviews and observations with philosophical conceptual analysis, we explore the uncertainties in this process of tool development at the intersection of scientific and clinical reasoning around disorders of consciousness.

We specifically trace how different kinds of uncertainties (arising from different sources such as probability, ambiguity and complexity) are salient in these different contexts (research and clinic) and how the specific relevance attached to these uncertainties in the context influence when and how to share results of individual tests produced in the process of research, and the role it plays in determining how to implement the new tools from research into clinical practice. Our findings suggest a clear difference in assumptions about how the immediate results of the experimental test should be used in the concrete assessment of the subjects participating in the project.

Our investigation has a descriptive focus, however uncovering these details may serve as a basis normative and ethical considerations in relation to selecting strategies for management of uncertainties, as well as in evaluation of evidence in research and the clinic in the future.

3. The Cincinnati Water Maze in the Making
Nina Atanasova, Charles Vorhees and Michael Williams

In this project, we adopt integrated methodology in presenting a case-study from experimental neuroscience. It exemplifies the interplay between theory, experiment, and technology. We show that, contrary to traditional accounts of science, tool-development and experiment rather that theory drive scientific change. Our collaboration aims at providing a comprehensive account of the invention and development of an experimental apparatus, the Cincinnati Water Maze (CWM), which was invented and has been continuously developed in the Vorhees/Williams Neurology Lab at Cincinnati Children's Research Foundation. In this paper, we detail the key steps in the development of the CWM. We trace the solutions to epistemic problems against the background of material and institutional constraints. We show that the invention and development of the CWM is a clear case in which tool-development advances independently from theory.

4. Interventionist Accounts of Causation Work in Experimental Practice
Tudor Baetu

The scientific defence of mixed-level causal models (e.g., biopsychosocial models of pain) rests on the notion that the elucidation of the causal structure of the world is a matter of experimental inquiry. If it is possible to intervene on a variable and measure a change in an outcome, then the variable is causally relevant to that outcome. It doesn't matter that the variable in question is psychological or social rather than biological. The experimentalist's argument should also extend to interventionist accounts of causation, including the highly influential account developed by James Woodward. Since the experimental scientist is free to test the causal relationships between any two variables, it seems only natural to conclude, as Woodward himself does, that in the context of an interventionist account of causation, "there is no bar in principle to mixing variables that are at what might seem to be different 'levels' in causal claims."

It turns out that this conclusion is problematic in at least one respect. Michael Baumgartner points out that, in order to provide conclusive evidence for causation, interventions must satisfy a fixability clause and the requirements of an ideal intervention. However, if higher-level variables supervene on lower-level ones, it is impossible to intervene on supervening states without also intervening on their supervenience bases. This creates a confounding problem, since it leaves open the possibility that changes in variables other than the intended target of the intervention may be responsible for (and therefore explain) the observed differences in outcome. The upshot of these limitations is that it is impossible to conclusively demonstrate that higher-level (e.g., psychosocial) variables are causally relevant.

Woodward's response to the challenge is that "it is inappropriate to control for supervenience bases in assessing the causal efficacy of supervening properties." He justifies this claim by arguing that matters of definition and metaphysical necessity act as prior constraints on possible causal relationships between variables. The proposed solution is therefore to exclude variables standing in supervenience, definitional, nomological, identity and other non-causal relationships with the tested (independent) variable as potential confounders. This effectively entails a modification of interventionist accounts in order to allow for 'fat-handed,' or inaccurate, interventions targeting simultaneously several variables standing in non-causal relationships.

In this paper, I argue that Woodward's solution constitutes an unacceptable compromise undermining the empirical foundations of science. My assessment is based on three considerations. First, no prior knowledge about the nature of confounders–be them causal, supervenience bases or something else–is needed in order to design and conduct an experiment testing the causal relevance of a variable. Second, the consensus in experimental science is that confounded studies cannot yield conclusive results because they fail to differentiate between rival explanations. And third, the concept of definitional and metaphysical constraints proposed by Woodward is ultimately based on a spurious example.

5. Revisiting the Distinction Between Predictions and Projections
Marina Baldissera, Bradley Seamus, Suraje Dessai and David Stainforth

Climate scientists distinguish between climate predictions and projections. This distinction intends to make an epistemic distinction between to different kinds of model based inferences in climate science. In this paper we argue that this distinction does achieve its intended aim, and that the distinction between projections and predictions, if one is to be made at all, is one of temporal scales of model output. One further worry with this distinction is that the words are used interchangeably by climate scientists, and the distinction drawn by the IPCC has been shown to be only loosely adopted when it is used at all.

Global Circulation Models (GCMs)—the models that provide a large part of the scientific information on which statements are based—are said to produce scenario dependent projections rather than predictions. Projections are defined as model output that is explicitly dependent on the assumptions about future scenarios that may or may not be realized. As a consequence, these projections are conditional statements about possible future climate conditions. Climate predictions, on the other hand, are defined as model output that attempt to represent the actual state of future climate, starting from particular initial conditions (see Annex II, WG II, AR5). As it stands, this distinction is ambiguous. On the one hand, this distinction is epistemic: predictions are about actual future states of the world, whereas projections are about possible states of the world, conditional on, for example, different greenhouse gas (GHG) emission scenarios. On the other hand, in so far as GCMs are descriptions of phenomena in the world, they all rely on our assumptions about how to represent these phenomena, and how the phenomena will evolve in the future. Hence, all GCMs are conditional on the assumptions about the world and how it evolves, and are, as a consequence, projections. It would be impossible, according to this interpretation, to obtain predictions from GCMs.

One way in which this distinction can be useful, however, concerns the temporal scales of the models and the relation between temporal scales and assumptions introduced in the models. For example, for short temporal scales, the so-called stationarity assumption is a reasonable assumption, as the changes in GHG occur on timescales that are much longer than the processes that are considered for short temporal scales. However, when models (and in particular GCMs) explore the climatic conditions for longer temporal scales, the stationarity assumption is no longer valid: more assumptions about what factors will affect these climatic conditions need to be introduced. For example, different levels of GHG will yield different climatic conditions – and these need to be taken into consideration. The relevant distinction between predictions and projections is therefore the extent to which assumptions used in model building determine what kind of inferences about the world can be made from models.

6. Identity by Descent (IBD) as an Algorithm-and-Simulation Practice in Current Human Population Genomics Research
Carlos Andres Barragan and James Griesemer

We present in-progress archival, ethnographical, and conceptual analyses of how life-scientists currently frame and implement the concept of identity by descent (IBD) as a model to predict biological events for "recent" or "regional" human ancestry. Although IBD can be contextualized to have multiple histories, here we focus on how the concept is used in biomedicine and population genomics to describe a DNA segment shared by two or more individuals inherited from a common ancestor. We argue that the use of IBD to substantiate relatedness among humans through time and space is less than a straightforward modeling and rendering practice following a well-described concept from population genetics. From a bioinformatics point of view, the use of DNA microarrays to measure genetic similarity involves a considerable risk of misinterpretation of false positives, e.g. due to recombination, complicating the understanding of "recent" and "deep" ancestry. A practice scientists seem increasingly to use in order to overcome such challenges is to devote more effort and resources toward building tailored analytical tools, such as algorithms and software packages, as means of exploring alternative modeling strategies. Framed by scientists as a key form of scientific innovation, these tailored tools are valued not only for helping researchers to address false positives, but more importantly, for allowing them to answer specific questions about contemporary human populations, e.g. individuals reproducing within relatively isolated communities. The landscape of datasets to which scientists seek to apply these tools is dynamic and ever-changing. New models developed through algorithm-and-simulation practices must be applied while relying on a palimpsest of datasets produced in various contexts (e.g. ancestry studies and biomedicine for new and different uses). We argue that IBD, now understood as a tool, can embody models, predictions and datasets as they are re-situated and re-used by multiple actors —e.g. researchers in particular labs, communities, individuals— engaged in tracking patterns of human ancestry and biomedical risk. In order to illustrate how scientific knowledge about human heredity and ancestry is re-situated among contexts we analyze several recent research outcomes from three university research laboratories in California focusing on human population genomics.

7. The Logic of Scientific Proxies

William Bausman

Scientists use proxies to learn about difficult to access parts of the world. Geologists use the relative ratios of gases trapped in bubbles in the ice as a proxy for atmospheric temperature based on physical and chemical processes of isotopes. Biologists use tree rings as proxies for past temperature and dryness based on ecological and physiological conditions favoring and disfavoring growth. Alfred Russel Wallace used the long nectary spur of an orchid from Madagascar as a proxy to infer the existence of a pollinator moth with a long proboscis that sucks it based on adaptation by natural selection. Philosophers have said little about the use of proxies in science and I aim to characterize the epistemic features of this powerful area of scientific reasoning. In this paper, using a contemporary case study of teeth as a climate proxy, I reconstruct how scientists establish and use proxies and relate this to two philosophical distinctions concerning the logics of discovery, pursuit, and justification.

The first distinction is between hypothesis formation and pursuit vs. hypothesis testing and acceptance. The debate between Peircian Abduction and Inference to the Best Explanation (IBE) concerns this distinction. Abduction for Peirce is the form of inference whereby scientists generate new hypotheses and come to pursue a hypothesis explanatory of a set of facts. IBE is the form of inference whereby scientists come to accept some hypothesis explanatory of a set of facts. The second distinction is whether evidence is gained indirectly or directly. The debate between IBE and Vera Causa standards of evidence concerns this distinction. In IBE, a hypothesis is evaluated on how good of an explanation it would be if it were true and the best one is accepted. This is an indirect, consequentialist type of reasoning. The Vera Causa ideal holds that scientists must show three things to infer a cause: the existence, the competence, and the actual responsibility of the cause. Most characteristically, the existence and competence of a cause must be established directly, independent of the consequences of its truth.

My case study shows how the research group led by evolutionary paleontologist Mikael Fortelius establishes and uses mammalian teeth as a proxy for past environmental chance based on adaptation. Relating their work to the logics of discovery and justification, I first show how their reasoning uses Peircian Abduction to generate pursuit-worthy hypotheses, not IBE. I then show how they combine Vera Causa with consequentialist reasoning. The scientists take pains to establish that the teeth they are using were adapted by natural selection to specific environmental conditions directly, without resorting to comparisons of predictions or theoretical virtues. Next, they use robustness reasoning to check the predictions of the teeth as a proxy with other, already established climate proxies. Since robustness reasoning is consequentialist, this case study shows that Vera Causa reasoning is complimentary to consequentialist kinds of reasoning within the logic of justification. I hope that this brings new attention to the kinds of methodological choices made by scientists and to the different activities that go on in the scientific practice across the contexts discovery, pursuit, and justification.

8. Modeling Practices in Ecology: Islands as Material Models

Saliha Bayir

Models have a central place in ecological research as they do in other biological sciences. Most ecologists make use of mathematical models and their representative function to generalize or extrapolate, especially for population-level phenomena like abundance, distribution, competition, etc. However, looking at the modeling practices of ecologists from history and philosophy of science in practice perspective shows that the material activities of ecologists are in most cases, intertwined with theoretical modeling (Griesemer, 1990) Even though material models seem scarce in ecology at first glance, they, too, are more commonly employed than it is acknowledged by ecologists. For instance, Griesemer highlights the importance of the remnant models of faunas (1990) and model taxa (2013) in ecology. However, philosophers are mainly interested in the role of mathematical models in ecology, although material models are proven to be fruitful for inquiry in other fields such as model organisms in molecular biology.

There are, of course, many other material models, especially when considering the range of experimental sites in ecological research. In this paper, I aim to investigate the modeling role of islands in ecology. Ever since MacArthur and Wilson's island biogeography model from 1960s (Odenbaugh, 2011; Warren et al., 2015), different islands have been seen as a convenient middle ground between tractability (of the experimental site) and complexity (of the ecological phenomena in larger continents): for example, in order to examine soil ecosystem nutrient fractioning, ecologists choose drier sites like Hawaiian Islands (Austin & Vitousek, 1998), to study tropical forest ecology, researchers mainly focus on the human-made Barro Colorado Island as a model (Leigh, 1999). On the basis of these two case studies, I will analyze the methodological and theoretical assumptions, the practices of data generation, use, and interpretation. This will provide an understanding of the ways in which islands are turned into exemplary experimental sites that have the capacity to represent a broader class of ecological systems despite, or rather because of their specific properties.

Austin, A. T., & Vitousek, P. M. 1998. Nutrient Dynamics on a Precipitation Gradient in Hawai'i. *Oecologia* 113(4).

Griesemer, J. (990. Modeling in the Museum: On the Role of Remnant Models in the work of Joseph Grinnell. *Biology and Philosophy* 5(1).

Griesemer, J. 2013. Integration of Approaches in David Wake's Model-taxon Research Platform for Evolutionary Morphology. *Studies in History and Philosophy of Biological and Biomedical Sciences*, 44(4PtA).

Leigh, E. G. 1999. *Tropical Forest Ecology: A View from Barro Colorado Island*. Oxford University Press.

Odenbaugh, J. 2011. Philosophical Themes in the Work of Robert H. Macarthur. In K. deLaplante, B. Brown, & K. A. Peacock (Eds.), *Philosophy of Ecology* (pp. 11–109). North-Holland.

Warren, B. H., et al. 2015. Islands as Model Systems in Ecology and Evolution: Prospects Fifty Years after MacArthur-Wilson. *Ecology Letters*, 18(2).

9. Placeholder Concepts in Investigative Practice: The Case of "Slow Infections"

Corinne Bloch-Mullins and Amesh Adalja

Recent work in the philosophy of scientific concepts has examined their productive roles in investigative practice (e.g., Feest & Steinle 2012). For example, rather than merely serving as vessels for previously accumulated information, concepts play a role in the empirical individuation of the phenomenon of interest, allowing for further generation of knowledge. This entails that concepts must be formed before scientists have a complete grasp on the entities in question (see, e.g., discussion in Feest 2010). This paper examines the generation and use of such "placeholder concepts" (Andersen 2012) in light of recent work on concepts and similarity in cognitive science.

Specifically, we ask: when does a new phenomenon warrant the formation of a new scientific concept, rather than an expansion of an existing one? Moreover, how does such a new category, which may lump together phenomena that turn out to be heterogenous at a more causally fundamental level of description, lead to the development of later, more homogenous categories? We address these questions by examining the formation and evolution of the concept "slow infection," which paved the way to the later concept "prion."

The category of "slow infections" was proposed in 1954 by Björn Sigurdsson. While Sigurdsson believed these diseases were caused by viruses, he suggested that they be distinguished from acute and chronic types of infections. Specifically, in forming this concept, he focused on attributes such as a long incubation period and a regular, predictable constellation of symptoms leading to death. These features had been known—and were associated with individual infectious diseases such as scrapie—long before Sigurdsson's proposal. What prompted him to suggest the formation of a new category?

Drawing on Rosch's work on similarity and conceptual taxonomy and on more recent work on the role of alignment in similarity comparisons and categorization, we discuss the interactive process that gives rise to the formation of a new category. Specifically, it is not enough that a phenomenon displays some anomalous features; rather, these features need to placed in correspondence with features of familiar phenomena that serve as contrast classes. The alignment, in turn, allows for further grasp of commonalities among the features of the members of the new categories.

Finally, we discuss the way in which the preliminary category of "slow infection" paved the way to narrower, more homogenic categories and the development of the "prion" concept. The paper thus draws on insight in cognitive science to shed light on the role of concepts in supporting their own evolution as part of a process of epistemic iteration (Chang 2004).

10. Molecular Models as Tools for Gaining Information

Agnes Bolinska

What makes model building useful? I approach this general philosophical question through a case from the history of scientific practice: How did building molecular models contribute to Linus Pauling's discovery of the alpha helix and Francis Crick and James Watson's determination of DNA structure? In an era when priority discovering molecular structure conferred considerable professional reward, effective heuristics were crucial. I argue that, by physically encoding information about bond types, lengths, and angles, molecular models further enhanced the efficiency of the most fruitful heuristic for molecular structure determination available to historical actors.

Bolinska (2018) construes the problems of macromolecule structure determination as beginning with a space of candidate structures, which can be successively reduced by considering various pieces of evidence serving as constraints on structure, each of which eliminates candidate structures. She argues that the most effective heuristic for determining molecular structure is one that reduces the size of the possibility space with the greatest degree of confidence that this has been done correctly. A components-first strategy for determining molecular structure, in which bond types, lengths, and angles are used to eliminate candidate structure before whole-molecule information derived from X-ray diffraction photographs, was likeliest to lead to the correct structure more efficiently than the photos-first strategy, which reverses this order, because it warranted greater confidence. Bolinska suggests that building molecular models was one way such a strategy can be applied.

In this paper, I extend this suggestion. I argue that physical molecular models acted as tools for gaining information about molecular structure by physically instantiating component-level constraints, providing a concrete means of prioritizing these constraints in scientists' reasoning and ensuring enforcement of stereochemical rules. For instance, sticks representing bond lengths were constructed to scale, so structures constructed from them were appropriately constrained by known bond lengths. Further, physical properties of the structures constructed from them precluded their violating certain stereochemical rules, which dictate how molecules may fold. For instance, a molecular model of DNA that pairs like bases with one another would be mechanically difficult to construct, since it would pinch in some places and bulge in others. And stereochemical considerations dictate that molecules adopt the least strained conformation possible—one that does not permit such pinching and bulging. Thus, molecular model-building enforces conformity with stereochemical rules.

Further, constructing molecular models serves as a cognitive aid, enabling the application of several component-level constraints simultaneously. The instantiation of some of these constraints by the model frees up some of a scientist's limited cognitive capacity for the consideration of other constraints. For instance, model-building permitted, among other things, the ready comparison of DNA structures that had like-with-like base pairing and those that did not. Thus, models even further reduced the size of the possibility space remaining upon the application of a given constraint.

My proposal is informed both by theoretical considerations about modeling and representation, and by how Pauling and Watson and Crick themselves conceived of model-building as a tool for finding molecular structure.

11. What are the 'Orchids' Sensitive To? Measuring the Environment in Differential Plasticity Research
Olesya Bondarenko

Gene-environment (GxE) interaction studies have been widely used in psychological, psychiatric and behavioural science to study outcomes ranging from depression to parenting behaviour. The idea behind the approach is that the outcomes of interest cannot be explained by the main effects of genes (or environments) alone but result from causal interaction between specific genotypes and environmental factors. As the GxE research has amassed vast amounts of empirical evidence, the scientists have faced the need to "stabilise the phenomena" (Feest 2009) from the data collected. What is the nature of the observed interaction and why do different genotypes generate such dissimilar outcomes in the presence of one and the same environmental factor? A number of scientists have posited the existence of a genetic vulnerability that increases the likelihood of a negative outcome under exposure to an adverse environment. However, an alternative way of parsing has been proposed by the developmental psychologist Jay Belsky and his collaborators (Belsky, Bakermans-Kranenburg & van Ijzendoorn 2007; Belsky et al. 2009; Belsky & Pluess 2009; Belsky & Pluess 2013; Belsky & Pluess 2016; Belsky & van Ijzendoorn 2017). They have argued that in a number of cases, the underlying phenomenon is not merely a vulnerability to adversity but "differential plasticity" with regard to environmental influences, both positive and negative. The research on "differential plasticity" has sought to demonstrate that some individuals are much more sensitive than others to environmental factors at both ends of the extreme (according to a metaphorical description used in the literature, such individuals are "orchids" rather than "dandelions").

In this paper, I examine both the theoretical framework of "differential plasticity" research and the existing empirical studies aimed at uncovering the purported plasticity effects with regard to several genetic variants, environments, and outcomes. I argue that despite potential epistemic merits of the underlying theory such as scope and external consistency with both evolutionary biology and a body of work in psychology and developmental science, much of the research conducted so far has left epistemic gaps stemming from 1) coarse-grained measures of environments and potentially omitted environmental variables, making it difficult to know which exact influences individuals are sensitive to 2) lack of precise and testable mechanistic hypotheses on the nature of the observed gene-environment interaction(s). In order to illustrate these points, I analyse a study (Beaver & Belsky 2012) on the role of differential genetic plasticity in cross-generational transmission of parenting. I conclude that increasing sophistication with regard to the genetic component (such as the use of composite measures of genetic plasticity) is not a suitable replacement for fine-grained and imaginative thinking about environmental influences, and suggest that the more recent turn to experimental studies may help rectify some of the mentioned shortcomings of "differential plasticity" research, provided that several conditions are met.

12. Philosophy of Science in Educational Practice: The Issue of Training Academic Skills for Complex Problem-Solving
Mieke Boon

Higher education policy documents promote the importance of academic skills, such as critical thinking, reflection, responsibility and interdisciplinarity, as crucial to the ability to deal with complex social challenges. Current educational practices train these skills through problem- or project-based-learning. Engineering education programs, for example, consist of thematic modules in which students are expected to conduct research for solving a real-world problem (e.g., the development of a biomaterial for a specific biomedical problem) by applying scientific knowledge gained in the associated courses.

Although these educational approaches work quite well for developing professional skills (e.g., communication, collaboration and organization), they are less successful for developing academic skills. Teachers experience that the level of academic skills remains low, with the weak points being 'the formulation of the problem,' 'the application of scientific knowledge and methods,' and 'critical and creative thinking.' Indeed, the educational literature shows, firstly, that this is a general phenomenon, and, secondly, that project-based education expects students to learn academic skills more or less naturally. In other words, there is no explicit training involved.

In some educational programs, attempts are made to address this issue through philosophy of science courses. However, philosophers who are involved in this type of service-teaching, and who expect (or are expected) to contribute to the academic skills of students, have often experienced that they are not as successful as they would like.

In this paper, I aim to address the issue of training the academic skills that are needed for doing research or taking scientific approaches in solving (complex) real-world problems, and to give a few suggestions about the possible roles of philosophers of science. Firstly, I aim to make plausible that science education in academic programs is still very much theory focused, thereby suggesting that solutions to real-world problems can be deduced from theories. This way of teaching science make it very hard to see how science is used in real problem-solving. I will hark back to Giere (1994), to defend that science teaching should not only focus on understanding the products of science (in particular, the internal structure of theories), but also at the research process in which the creative contribution of the (creative and critical) human cognitive system and of human interactions through experiments and technological instruments become clear. Secondly, I will explain a new educational vision that shifts the focus from theories to problem-solving through conceptual modelling, which comprises four levels:

- Philosophically, it shifts the emphasis from knowledge as a product to the role of the research process (as an alternative to positivism and relativism).
- Epistemologically/methodologically, it emphasizes the role of academic skills and technology in knowledge formation (as an alternative to focus on 'true' knowledge).
- Normatively, it addresses the crucial role of epistemological responsibility in research.
- Practically, conceptual modeling combines teaching a research method and training academic skills.

Giere, R.N. 1994. The Cognitive Structure of Scientific Theories. *Philosophy of Science*, 61(2).

13. What's at Risk in Scientific Risk?

Marion Boulicault

The value-free ideal (VFI) of science has come under criticism over the last 70 years from what I call risk arguments. These take the following form: core scientific reasoning requires risk management; risk management requires values; so, core scientific reasoning requires values. An early formulation of the risk argument was restricted to a particular kind of risk (inductive risk) that occurs at a particular stage of scientific practice (setting certainty thresholds for hypothesis acceptance) (Rudner 1953). Defenders of the VFI countered that the setting of certainty thresholds doesn't involv--e 'core' scientific reasoning: the VFI is safe (Jeffrey 1956). In response, advocates of risk arguments have adopted what I call the risk proliferation strategy: many kinds of risk occur (Biddle and Kukla 2017) at many stages of scientific practice—including the generation of evidence itself (Douglas 2000, Kukla 2017). According to the risk proliferation strategy, risks reach the core of science.

Expanding on arguments by Brown (2013), Brigandt (2015) and de Melo-Martín & Intemann (2016), I show that the risk proliferation strategy is doomed to failure. Common to all risk arguments is the notion that the risk incurred when making scientific choices is the risk of "getting it wrong": when we choose a significance level or define a disease, we risk making a choice that gets it wrong, e.g. accepting a false hypothesis, or providing unnecessary medical treatment (Biddle 2016, 202). Values play a role only in determining how to manage risks of error, but not in determining what counts as an error in the first place, or when one has occurred. As such, the risk argument leaves a core part of scientific reasoning insulated from values. Risk arguments implicitly endorse a value-free core of science in aiming to show that, even in its core, science isn't value-free. This risk proliferation strategy is thus the wrong strategy: no matter how many kinds of risk we postulate, or how deep these risks run, risk arguments can never defeat the VFI.

I illustrate my argument by way of a case study: a recent risk argument (from Kukla 2017) concerning evidence generation via infertility measurement. I unearth a contradiction in Kukla's argument. Like other risk theorists, Kukla explicitly rejects the VFI, arguing that values are necessary, even for measurement. Kukla uses her risk argument to argue that infertility measures are problematic and should be rejected, but, crucially, she does so by arguing that the values in infertility metrics are "buried too deep" such that infertility cannot be conceptualized independently of "value-laden social factors" (2017, pp. 4426, 4419). In other words, Kukla simultaneously accepts that measurement is value-laden while rejecting infertility metrics because they are too value-laden! I demonstrate how this contradiction arises from Kukla's use of risk: she rejects the VFI while using an argument that implicitly endorses it.

I end by drawing on lessons from this case study to suggest more promising approaches for conceptualizing the role of values in science, approaches that leave risk behind and successfully go beyond the VFI.

14. Stars in Jars and Hybrid Black Holes: Interpreting Empirical Results in Astrophysics without the Observation/Intervention Dichotomy
Nora Boyd

The traditional distinction between observation and intervention is ill-suited to rich philosophical investigation of science in practice. It is perhaps tempting to define empirical science as investigation of the natural world via the methods of observation and targeted intervention in controlled experiments. Yet, I will argue that observation and intervention are irrelevant to the empirical nature of scientific evidence, and hence for its epistemic utility in learning about nature. Instead, what matters is that the evidence derives from a causal chain that has one end anchored in the worldly target of interest. I show that invoking the observation/intervention dichotomy, and failing to track the distinction between the empirical and the virtual, leads to three mistaken inferences. I then argue that attending to the distinction between the empirical and the virtual yields deeper and more interesting questions about the epistemology and methodology of science, drawing on illustrative cases from contemporary astrophysics.

Philosophers have claimed that astrophysics is a distinctively observational science (cf. Anderl 2016), that it invites the use of simulated data in place of experimental data (cf. Morrison 2015), and even that it is not a properly scientific field at all (cf. Hacking 1989). These three theses all rely on a fraught observation/intervention dichotomy, and attending to astrophysics in practice yields many cases that cross-cut or side-step this distinction.

Researchers at the National Ignition Facility are attempting to study astrophysical phenomena, like hydrodynamical instabilities in supernovae, with high-energy-density states of matter created by irradiating small pellets of plastic and foam with powerful lasers in a laboratory setting. The epistemic success of this endeavor does not depend on whether this research is properly construed as “observation” or “intervention”. Rather, it depends on the researchers’ ability to make physically grounded similarity arguments that bridge the lab and the stars. In one recent example (Kuranz et al. 2018), it seems that researchers have unfortunately undermined premises needed for the for the relevant similarity argument via the very conditions of their research focus.

It is often in virtue of the transformation of data by significant processing that empirical results serve as constraints on theorizing at all. In the case of the Event Horizon Telescope image of M87*’s accretion disk, in order to check the consistency of the data with general relativity, the data were transformed, but they were also supplemented by modeling assumptions that helped to fill in details that made constructing the full image from limited sky coverage possible. Again, whether this research is an instance of “observation” or “intervention” proves unhelpful to understanding the epistemic significance of its results. Instead, their significance depends on the manner and extent to which synthetic data can be merged with empirical without breaking the key causal connection between the results and the worldly target. I argue that epistemic successes and limitations can be more readily revealed to philosophical scrutiny when we attend to the conditions that make scientific evidence empirical.

15. Toward a More Fine-Grained and Diverse Methodological Landscape of Human Behavior Research

Ingo Brigandt

Helen Longino has prominently argued that different approaches to the study of human behavior cannot be integrated. She even maintains that these approaches are incommensurable, among other things because they parse the overall space of causal factors in a disparate fashion. While this gets at the importance of methodological differences, I argue that Longino neglects overlap among and the potential for integration between approaches. Articulating human behavior research in terms of incommensurable, monolithic approaches also ignores the methodological diversity within a research field. By considering for instance molecular behavioral genetics as one ‘approach,’ Longino assumes that across a whole field one overarching methodology is being adopted. This talk, in contrast, works toward articulating a more fine-grained landscape of human behavior research. I investigate research in terms of individual methodological strategies and explanatory commitments, where a research group adopts several such commitments.

This is illustrated by an examination of epigenetics research on neuropsychiatric conditions. I argue that the primarily molecular approach of epigenetics exhibits heterogeneity, where some (though not all) methodological commitments found in the overall field have the potential for coordinating between molecular epigenetics research and approaches that investigate social influences on neuropsychiatric conditions, e.g., psychiatric epidemiology and social psychiatry. This possibility for integration is something that Longino denies, when maintaining that molecular behavioral genetics and social-environmental approaches are two incommensurable approaches.

Epigenetic modifications impact gene regulation and are involved in neurocognitive processes. As epigenetic modifications are set up and modified in response to a cell's or person's environment, this has motivated research on neuroplasticity in humans (including in psychiatric contexts), which promises a mechanistic account of gene-environment interactions. At the same time, some epigenetics researchers addresses epigenetic marks that are set up during pregnancy or early childhood. When these are taken to have a lasting impact on symptom formation, this amounts to a more deterministic vision that fails to investigate how the social environment impacts psychiatric symptom formation throughout a person's lifetime. Epigenetics research on neuropsychiatric conditions can also reveal new behavioral and psychotherapeutic treatments, but this does not hold for research groups focusing on potential pharmacological treatments that would biochemically modify epigenetic marks. There are gender differences in many neuropsychiatric conditions. Some epigenetics research is open to investigating the full range of between-person variation in psychiatric features and social-behavioral influences, while other studies focus on contrasting a female and male condition or attempt to reveal some biological basis of sex-differences, so as to factor out social influences on gender differences in neuropsychiatry.

Longino is right to focus not on theories but on methodology (including conceptual aspects, e.g., parsing causal spaces), because methodological features have implications for the future, including forming a hindrance to integration. At the same time, my fine-grained perspective reveals the presence of alternative methodological strategies among some molecular scientists, which motivate integration or hold the potential for integration with social-environmental research.

16. Evaluating Metrics for Interdisciplinary Integration

Evelyn Brister and Marisa Rinkus

Interdisciplinary research faces unique challenges when it addresses research problems that require integrating the social and natural sciences. Differences between sciences include not only different sets of factual and conceptual knowledge, but also differences in methods, standards of evidence, causes, and values. This latter set of differences presents a potential obstacle to producing new, integrated knowledge because judgments concerning methods, rigor, causal interactions, and values are fundamental to evaluating research quality. Thus, these kinds of differences sometimes lead to misunderstandings and disagreements within a research team that are more serious than simple miscommunication. They are more serious because there is no clear external, rational standard that can resolve them. At the same time, carefully working through such epistemic differences can be essential to integrating different forms of scientific knowledge and producing innovative problem-solving approaches.

Our goal in this paper is to evaluate how measures of interdisciplinary integration address negotiating epistemic obstacles between disciplines. We are motivated by the imperative to avoid integration failure: it is unfortunately rather common for social scientists to report that they were added too late to an interdisciplinary project to integrate their insights into the research design.

There are several ways that interdisciplinary integration is conceptualized, but the most common measures of integration do not capture how social scientists report integration failure. Instead, the most common metric of interdisciplinary integration, the Rao-Stirling Integration Score, measures, for a corpus of research articles, how widely and how far from its locus the article set cites literature across a conceptual map of disciplines. Such metrics conceptualize interdisciplinary integration in terms of the connections that research articles make between bits of knowledge. Interdisciplinary connections are demonstrated through their citation networks and are thus public and relatively easy to measure.

Such metrics focus on the content of research products. However, such bibliometric integration scores are unable to identify the failures described by social scientists because they assess knowledge content rather than the success or failure of the process of interdisciplinary collaboration. For the most part, failures to integrate knowledge to the satisfaction of social scientists remain invisible. Research projects may produce results, and they may even produce results that cite widely across disciplines, but they may nonetheless produce results that lean heavily to the epistemic presumptions of a single discipline rather than integrating knowledge across disciplines.

We propose that evaluating the kind of knowledge integration that avoids disciplinary capture requires assessing the collaborative process and division of labor for research teams. We examine empirical data that distinguishes interdisciplinary teams employing a process where intellectual labor is evenly distributed from teams employing a parceled division of labor. We argue that a more even distribution of labor provides more opportunities to work through epistemic assumptions about methods, standards of evidence, causes, and values.

17. Inferentialism and Maker's Knowledge

Daniel Burnston

There are two broad approaches to scientific representation, and how it relates to explanation. The first is what I call "referentialism," which locates the representational and explanatory relationship between a model and the world in some correspondence relation between them. The second approach is "inferentialism," which argues that the representational and explanatory power of a model lies in the relation it has to scientists – namely the way it affords and constrains inferences about the world.

At first glance, referentialism seems to have clear epistemic advantages over inferentialism, since explanatory success is closely tied to referential success. Hence, a successful explanation bears a very close relationship to the world. On the other hand, inferentialism seems inextricably tied to human reasoning processes, and hence risks being psychologistic – i.e., granting explanatory success to models without guaranteeing any objective relation to the world we are trying to explain.

My moves in this paper are two-fold. First, I argue that the apparent advantage of referentialism in explanatory objectivity is a false one. Referentialist views, when faced with widespread abstraction and idealization in science, trend strongly towards fictionalism – the view that models explain something, but that thing isn't the world. Inferentialism, since it does not base explanatory power on structural correspondence, is in no such situation.

Second, I then propose a positive epistemology for inferentialism based on the notion of maker's knowledge. I briefly consider and reject two accounts of maker's knowledge, the "recipe model" and the "looping model," since they fail to solve the epistemic problems for inferentialism. On the recipe model, maker's knowledge is due to having the ability to literally create the thing that is understood. But this fares no better than referentialism at generating knowledge of the actual world. On the looping model, one's actions shape the world into being a certain way (hence, looping is often discussed in the context of social images like pornography, and diagnostic categories such as mental disorders). However, this is not an adequate account of what scientists do, since modelers do not literally create the real-world systems they understand through their models.

As an alternative to the recipe and looping views, I propose an "engineering model" of maker's knowledge, on which practitioners know how to construct artifacts (scientific representations), and learn about the world through what those artifacts allow them to do. I show how, as with other artifacts, scientific models are successful or unsuccessful depending on how well they fulfill their purposes. Unlike many other artifacts, however, the purpose of models is epistemic. A model succeeds or fails depending on whether it can be used to generate confirmable predictions about real-world system. In sum, scientific models are artifacts and scientific modelers are conceptual engineers. I support my claims with examples from systems biology and systems neuroscience.

18. The Central Dogma Is Not Useful to Biology
Marco Camacho

The Central Dogma—which roughly says that DNA makes protein and not the other way around—has been a highly-influential principle in biology. The Dogma has been viewed as a catalyst for the Modern Synthesis in biology (Roll-Hansen 2011), as definitive evidence against Lamarckian evolution (Maynard Smith 1993; Judson 1979; Dawkins 1970; Cobb 2017; Wilkins 2002; Graur 2019), and as evidence supporting the gene's eye view of the development and evolution of an organism's phenotype (Rosenberg 2006; Weber 2006). Some have even claimed that the Central Dogma can be justified using computational methods (Lin and Elowitz 2016).

What's more, since the Dogma's formulation by Francis Crick in 1958, many commentators have debated the Dogma's empirical adequacy (see e.g. Camacho 2019; Crick 1958 & 1970; Keyes 1999; Stotz 2006; Rosenberg 2006; Weber 2006; Griffiths and Stotz 2013; Griffiths 2017; Sustar 2007). In these debates, however, little attention has been paid to the Dogma's practical significance to biology, which is surprising since there are numerous ways in which a scientific principle may be of use to science even if it fails to be empirically adequate.

Most obviously perhaps, a—false—scientific principle may be of use to science as an idealization (McMullin 1985). A scientific principle may also engender an understanding of nature: as Morrison (2015) has argued, empirically inaccurate models—such as the Hardy-Weinberg Law—afford a better understanding of the "genetical basis of evolutionary change" than more empirically-accurate models (Morrison 2015, p. 33). Nancy Cartwright (1983) has also noted that certain laws—though false descriptions of what really happens—may aid in the construction of highly predictive models. Additionally, Jonathan Birch (2014) has pointed out that scientific principles, such as Hamilton's Rule in sociobiology, may posit a set of general features that help unify seemingly disparate phenomena under a single conceptual scheme. All of these examples suggest that the Central Dogma may be plausible on practical grounds alone, and yet fail to be empirically adequate.

My aim is to move beyond discussions of descriptive accuracy, and instead to evaluate the Central Dogma's significance to biology from a purely practical perspective. To do this, I consider four distinct approaches for determining the non-descriptive methodological and practical significance of a scientific principle (McMullin 1985; Cartwright 1983; Kitcher 1989; Elgin 2014), and determine whether the Dogma is made plausible by any of these approaches. In doing this, my hope is to consider the Central Dogma in an entirely new light. I conclude that none of these approaches vindicate the Central Dogma. What's more, under many of these approaches, the Dogma amounts to a triviality.

19. Is Helen Longino even an Empiricist? Understanding Contextual Empiricism within Feminist Empiricism
Christopher Choglueck and Elisabeth Lloyd

While philosophers of science have made strides in understanding the role of values in the sciences, there remain open debates about the normative relation between values, evidence, and practice. Thirty years ago, Helen Longino developed contextual empiricism in Science as Social Knowledge (1990), based primarily on knowledge from the sciences, yet still critical toward how dominant values like sexism shape scientific judgment. Longino has received sustained criticism: on one hand, some philosophers contend that her contextual empiricism is too inclusive of oppressive values to be feminist. On the other hand, some feminist empiricists criticize Longino for not holding values to the same level of empirical scrutiny as other claims in science. What should we make of these criticisms, particularly given the 30-year anniversary of Science as Social Knowledge? Is Longino's synthesis of feminism and empiricism misguided or hopeless? Moreover, how should we understand the relation between feminist empiricism and Longino's contextual empiricism?

This talk focuses primarily on critiques from Clough (1998), Goldenberg (2015), and Solomon (2012) that contextual empiricism is not sufficiently empiricist in its treatment of values. We begin by describing Longino's program, particularly her views on background assumptions and evidential status. We then focus on two lines of empiricist criticism: First, we discuss how others claim that Longino has a non-empirical understanding of values. They advocate a supposedly alternative approach of treating all "values as evidence." Second, we discuss how Longino's critics claim that she cannot provide any empirical support for feminist values—or empirical refutation of sexist values—within her contextual empiricism.

In response to these critiques from other feminist empiricists, we defend Longino's account as a strong form of feminist empiricism based on scientific practices. We argue that these criticisms are based on mischaracterizations of Longino's position and overstatements of certain claims. In addition, we contend that contextual empiricism not only allows for the empirical support/refutation of values, but Longino explicitly discusses when values can be empirically adjudicated and when not.

The primary significance of this talk is resuscitating contextual empiricism within feminist empiricism as a robust system for understanding the normative relation between values and evidence. In contrast to recent work on “values as evidence,” we show the limitations of empirical tests of values. In line with Longino’s work, we elaborate on conditions in which values can be refuted/confirmed empirically and conditions when such attempts are merely partial or completely ineffective. Accordingly, we demonstrate the prospects for a more tempered account of normative power of science regarding values.

Clough, Sharyn. 1998. A Hasty Retreat from Evidence: The Recalcitrance of Relativism in Feminist Epistemology. *Hypatia*, 13(4).

Goldenberg, Maya J. 2015. How Can Feminist Theories of Evidence Assist Clinical Reasoning and Decision-Making? *Social Epistemology*, 29(1).

Longino, Helen E. 1990. *Science as Social Knowledge*. Princeton: Princeton University Press.

Solomon, Miriam. 2012. The Web of Valief: An Assessment of Feminist Radical Empiricism. In *Out from the Shadows: Analytical Feminist Contributions to Traditional Philosophy*, edited by Sharon L. Crasnow and Anita M. Superson. Oxford University Press.

20. Believe Me, I Can Explain! Concerns about Inferences to the Explanandum
David Colaço

Philosophical accounts of scientific explanation typically identify the targets to be explained (explananda) and the explanations for them (explanantia). There is also analysis of explanatory reasoning, which includes but is not limited to analyses of inferences to the best explanation. However, what is less discussed is the defensibility of an inference in the other direction: should we infer to an explanandum from plausibly explaining it?

While this inference is seldom discussed by philosophers, it is no less pressing in light of psychological research. Consider explanation effects, where individuals express higher likelihoods in explananda when provided explanations for them. A prima facie example of this effect is the “soy boy” case. In this case, the explanandum is that the consumption of soy has “feminizing effects on men” (Messina 2010, 2095). Adherents claim that they believe in the soy boy effect in part because they can “explain” it: soy contains phytoestrogens, the ingestion of which causes changes in sex characteristics. This explanation has the proper form for one account of explanation: it sketches a mechanistic relation between soy consumption and bodily changes. The explanation also has some degree of empirical support: evidence corroborates that soy contains chemicals called ‘phytoestrogens’ and that hormonal estrogen can induce these changes in humans. However, there is no evidence that soy has this effect in humans.

This case of explaining as a means of supporting an explanandum’s adequacy, whether involving a good faith inference or not, seems to have considerable rhetorical power. This raises a question. In inquiry about science, when provided a plausible explanation for an explanandum, should we infer that this explanandum is adequate? Should we (say) believe that the soy boy effect occurs based on its empirically-supported, mechanistic explanation?

This paper critically addresses what I call inferences to the explanandum: inferences from the premise that an explanandum is plausibly explained to the conclusion that this explanandum (or a claim or representation thereof) is adequate or true. These inferences thus consist in inferring from the premise that one can plausibly explain why or how something occurs to the conclusion that it occurs. Psychological research and historical analysis reveal that inferences to the explanandum are proposed by lay people and scientists, and prominent philosophical accounts of explanatory reasoning permit these inferences. I argue that we should be skeptical of inferences to the explanandum, though distinct inferences to the explainability (and pursuit worthiness) of explananda are acceptable. My conclusions raise concerns for how both scientists and lay people assess the value of explaining in their reasoning practices. These concerns include the negative consequences of the compelling nature of explanations as well as the rhetorical strength explaining can play in epistemically-suspect belief formation.

Messina, M. 2010. Soybean Isoflavone Exposure does not have Feminizing effects on Men: A Critical Examination of the Clinical Evidence. *Fertility and Sterility*, 93(7).

21. Reciprocal Causation across Biology and Cognitive Science
Amanda Corris

In biology, the concept of reciprocal causation suggests that evolutionary processes are fundamentally intertwined with developmental processes. Arguments for the bidirectionality of developmental and evolutionary causal influences are seen as rejecting the distinction between proximate and ultimate causes. On this view, the individual organism plays an important role in both its own development and its species' evolutionary trajectory. Calls for an extended interpretation of contemporary evolutionary theory, such as those put forth by the proponents of the Extended Evolutionary Synthesis research program, emphasize the relevance of a notion of reciprocal causation. In contrast, received views of evolutionary theory are seen as prioritizing a level of explanation that arguably excludes relevant ontogenetic factors.

A parallel debate can be found in cognitive science. Embodied approaches to cognition emphasize the dynamic relationship between organisms and their environments, stressing that information flows bidirectionally and incorporating causal factors across several scales of biological organization, namely at the scale of the body. Cognitivism and related views suggest that cognition can be understood locally – as operating on a sublevel distinct from other levels of biological organization within the organism. On this view, information flows unidirectionally and in a feed-forward system. It is received as sensory input, processed according to a set of predetermined rules, and effectively transcribed into behavioral responses.

In this talk, I suggest that the concept of reciprocal causation can have a valuable epistemological and methodological impact on both biology and cognitive science. At the heart of the debates in each field is a defining claim regarding causal directionality, and drawing attention to the claims made by each side reveals similarities in underlying commitments that impact the conceptual architectures of both fields. Taking behavioral development as an example, I suggest that a bidirectional hypothesis motivates a holistic explanatory account of behavioral traits that integrates multiple scales of biological organization. While this account suggests a more complex picture of behavior, it also provides a way to synthesize findings across several cognate fields – namely developmental biology, molecular biology, and psychology – in turn potentially filling in gaps in scientific knowledge across multiple research programs. It can thus do significant work toward shared epistemic goals of biology and cognitive science more generally, namely explaining the mechanisms for species-typical behaviors, but also potentially contributing toward an understanding of the generation of novel behaviors. In both theory and practice, taking seriously the concept of reciprocal causation can result in a rich interdisciplinary framework for understanding behavior.

22. Towards A Functional Account of Explainable AI
Kathleen Creel

Despite advances in prediction and classification, opaque machine learning algorithms often do not provide explanations for their decisions. Much of the sizeable literature addressing this problem in the field of explainable AI focuses on the production of token lay explanations of opaque algorithms for specific purposes. Such applied strategies, while valuable, typically lack a theoretical analysis of what an explanation is and why the explanation produced should be considered adequate. They are also often divorced from an epistemic, truth-centered notion of explanation, using self-reported understanding or end-user satisfaction as the metric of success.

I present a pragmatist analysis of the goals of explainable AI and a functionalist framework in which explanatory goals can be matched to forms of explanation which may achieve them. Many forms of explanation are extant in the field: token explanations for singular cases, mechanistic explanation, and counterfactuals which might form the basis of interventions. However, it is unclear how to choose between forms of explanation. When is one kind of explanation rather than another appropriate? What criteria should we use to judge the appropriateness of the explanation for the context, or the context for the explanation?

Work within explainable AI suggests that the primary factor in determining the appropriate form of explanation is the role of the explainee. Although a role based analysis is important, it will not always be the right lens by which to determine purposes. Drawing on themes from the American pragmatists, especially John Dewey, and from logical empiricists such as Otto Neurath as well as from contemporary pluralist and functional accounts of scientific explanation (Mitchell 1997; 2002; 2019; Woody 2015; Danks 2015), I propose going beyond "it depends" to give an account of what it depends on. I will investigate the role that goals, capabilities, and environments play in fitting explanations to purposes.

Many with pragmatist inclinations argue for an 'epistemology as a set of hypothetical imperatives' in which rationality is properly aimed at determining the means to a predetermined end (Schulte 1999). Following Dewey, however, I will argue that the value of an end can only be determined once the feasible means to achieve it are themselves known (Dewey 2008, 13:210–19). The determination of a set of possible means towards an end and the weighting of those means against one another shapes the subsequent evaluation of the end. Ends and their related means are chosen together, as a package.

The reciprocal determination of means and ends suggests that explainees cannot be statically assigned a type of explanation based on their role with respect to the system.

Their ends, and therefore their roles, will change as they come to understand the system better. Explanation must allow not merely for the fulfillment of an epistemic goal but also for the determination of whether the goal is valuable and whether the proposed consequences are worth the effort.

23. Risk Factors versus Causes; Philosophy versus Epidemiology

Leen De Vreese

Although talk and thought in terms of risk factors currently abounds in healthcare - as if it has never been different - the use of the notion "risk factor" is a recent phenomenon. The risk factor approach emerged not earlier than in the 1950s and 1960s, when the notion of probability came into use in the epidemiology of chronic disease. In the western world, most infectious diseases where by that time under control, and the main disease burden came now from people suffering from chronic diseases. But while the infectious diseases could be approached as resulting from a single main external cause (the invading microbe which is a necessary cause for the disease to develop), the etiology (and, by consequence, prevention) of chronic diseases turned out to be much more complex and to require a different causal approach. This resulted in a shift in healthcare research from a monocausal approach to a multicausal approach. And meanwhile, from thinking in deterministic terms (at the individual level) to thinking in probabilistic terms (at the population level). During the second half of the 20th century and the first decade of the 21st century, talk in terms of risk and risk factors has become ever more pervasive. Skolbekken (1995) has shown that there was an explosive increase in the number of articles in medical journals containing the word "risk" in the title or abstract between 1967 and 1991. His findings led him to speak about a "risk epidemic" in the medical sciences. However, despite the wide use of the notion risk factor, the literature shows that there is no consensus about the use of the term. Foremost, it is unclear how the notion "risk factor" differs from the notion of "cause". Clarifying the difference between the two notions is the focus of this paper.

For our analysis, we first look at what different authors from the biomedical field have to say about the causal status of risk factors. We will discover that three different views about the relation between 'risk factor' and 'cause' can be discerned in the literature. To evaluate these different points of view, we will then first present a philosophical typology of type level causes, that will help us in our further analysis. Afterwards, we will look closer at what epidemiologists are actually looking for when looking for risk factors, how they discern causes from risk factors, and what kind of risk factors they discern in their theoretical writings (in terms of typologies of risk factors). We will bring all this together in our analysis in which we compare the philosophers' point of view with the epidemiologists' point of view. This leads us to argue, first, for an approach in terms of 'causal risk factors' versus 'non-causal risk factors' (rather than in terms of 'causes' (generally) versus 'risk factors' (generally)). Further, a more fine-grained distinction will be made among five possible meanings of the term risk factor that can be discerned on the basis of our analysis.

24. The Practical Foundations of Ethology: How Karl von Frisch's Experiments on Animal Sensation Reconciled Holism and Reductionism to Become an Exemplar for Classical Ethology

Kelle Dhein

In 1973, the discipline of ethology came into its own when three of its most prominent practitioners—Konrad Lorenz, Niko Tinbergen, and Karl von Frisch— jointly received the Nobel Prize for Physiology or Medicine. Historians have shown how Lorenz and Tinbergen were central to the practical and theoretical innovations that came to define ethology as a distinct form of animal behavior research in the 20th century (Heinroth and Burghardt 1977; Burkhardt 1981, 1983, 1997, 1999, 2005; Jamieson and Bekoff 1992; Brigandt 2003, 2005; Kruuk 2003; Beale 2008). Von Frisch is rarely mentioned in such histories (Jaynes 1969; Beer 1975; Heinroth and Burghardt 1977; Lea 1984, chapter 1), and when he is, he is often framed as an expert practitioner of ethology rather than a founding influence (Bates 1953; Thorpe 1979; Dewsbury 1992; Jamieson and Bekoff 1992, p. 111–112; Kruuk 2003; Burkhardt 2005, p. 6; Greenberg 2012; Radhakrishna 2018). Contrary to such narratives, this paper argues that Von Frisch's experimental approach to animal sensation functioned as an exemplar, in the discipline-building sense articulated by Woody (2003), for Lorenz and Tinbergen's early vision of ethology. To see how, one needs to understand what distinguished Von Frisch's experimental practices from those of his contemporaries and how those distinguishing features exemplified the disciplinary space Tinbergen and Lorenz sought to claim for ethology during the discipline's early formation in the 1930s and 1940s. During that period, Tinbergen and Lorenz sought to characterize ethology as a discipline that achieved a genuinely biological approach to the study of animal behavior through a reconciliation of reductionist and holist commitments. On the one hand, Lorenz and Tinbergen championed reductive causal explanations of behavior in the sense that behavior ought to be explained in terms of its underlying physiology rather than in terms of higher-level psychological states. On the other hand, Tinbergen and Lorenz promoted a holistic conception of the animal as a complex collection of interdependent parts whose behaviors bore nuanced relationships to their environmental context and evolutionary history. Due to those nuanced relationships, behaviors possess adaptive significance to the animals who perform them, and Tinbergen and Lorenz were keen to incorporate an awareness of that significance into behavioral experiments.

Von Frisch's experimental style became an exemplar for classical ethology because of the way it reconciled these reductionist and holist commitments. To show how, I compare Von Frisch's (1913, 1914) investigations into animal color vision during the early 20th century with those of his contemporaries. Doing so highlights features that distinguished Von Frisch's investigations within their historical context. There are a plurality of ways for scientific investigations to be more or less "reductive" or "holistic". By using the animal color vision experiments of ophthalmologist Carl von Hess (1909) and comparative psychologist Charles Henry Turner (1910) as a historically relevant comparison class, I make practice-based claims about the technologies and methods that distinguished Von Frisch's investigations into animal sensation as being more or less reductive or holistic. The Von Frisch-Hess comparison highlights Von Frisch's evolutionary thinking, naturalistic methods, and holistic concern for the way environmental context affects behavior. The Von Frisch-Turner comparison highlights Von Frisch's high level of experimental control, commitment to reductionist explanations, and unwillingness to make claims about the mental life of animals. In the end, this paper shows how the concrete practices and technologies of an experimental researcher were incorporated into the more abstract programmatic ambitions of biologist discipline-builders.

25. How to Make Value-driven Science for Climate Policy More Ethical

Justin Donhauser

In several papers, I explain how Probabilistic Weather Event Attribution studies (PEAs) work, and explore various ways they can be used to aid in UNFCCC climate policy decisions and decision-making about climate justice issues (e.g. how to effectively deal with climate refugee cases). In this paper, I critically evaluate the limitations of PEAs and consider the extent to which it is ethical for PEA researchers to attribute responsibility to specific countries for specific extreme weather events and to make claims about liability for losses and damages due to those events as they do. I focus on showing that PEAs routinely presuppose that the methods they use are not prone to inductive risks and presuppose that PEA researchers thus have no epistemic consequences or responsibilities. I will argue that although PEAs are nevertheless crucially useful for practical decision-making, the “attributions” of liability made by PEA researchers are in fact prone to indicative risks and are driven by value judgments that PEA researchers should make transparent to make such studies and attributions more ethical. The paper concludes with a discussion of the implications of my argument for the ongoing, and very influential, debate about how PEAs should guide climate policy and relevant legal decisions, and critical reflection on some related implications for Evelyn Fox Keller’s argument in “Climate science, truth, and democracy” (recently published in Studies C).

26. Disciplinary Shifts in Machine Learning

Ravit Dotan

This paper develops a conceptual framework to analyze a central kind of process in the discipline of machine learning: the rise and fall of families of algorithms, i.e. of model-types. The paper then uses this conceptual framework to show that social values play an essential role in evaluating model-types and illustrates using the currently predominant model-type, deep learning.

It is commonly thought that the rise of deep learning is a case of clear and “objective" progress in machine learning. We show why this account fails, and more generally why the comparison between model-types is not “objective". We argue that the comparison between model-types is model-type-laden, meaning that the comparison between model-types depends on considerations on which people who are committed to a different model-type would disagree upon. Further, we argue that the considerations to use in comparing model-types are also value-laden. In particular, we argue that considerations that favor deep-learning over its competitors implicitly encode social and political values around centralization of power, political freedom, and deprioritization of environmental impacts. In doing so, this paper contributes to conversations on the ethics of in machine learning. So far, much of the attention has been given to fairness issues in the context of particular algorithms. The analysis in this paper shows how fairness issues are integrated in big-picture process in the discipline at large.

Our critique is based on our analysis of the concept of a model-type. We argue that researchers in the field of machine learning are committed to model-types, and that this commitment shapes their research activities in several ways: it shapes the selection of problems; constrains the solutions that are considered; and promotes and reinforces certain prerequisites. For example, a commitment to deep learning involves: prioritizing problems in fields such as computer vision; presupposing that models should be evaluated in environments that are rich in data and compute power; and prioritizing predictive accuracy over all other virtues. We don't dispute that deep learning models do better than models of other types in these conditions. However, these presuppositions are not shared by those committed to other model-types. Therefore, arguing that deep learning is superior by evaluating deep learning models in environments that satisfy these conditions means stacking the cards in favor of deep learning. That is why the evaluation of model-types is model-type-laden. Moreover, the comparison between model-types is value-laden when the considerations that favor one model-type over another are value-laden. The considerations presently used to compare model-type are indeed value-laden. For example, presupposing that models should be evaluated in environments rich in data and compute power promotes the centralization of power. The need for specialized and expansive hardware as well as huge data sets creates an entry barrier that favors those with financial means, such as big companies in rich countries.

27. Conceptualizing the Public in Science Engagement Practices
Heather Douglas and Jared Talley

There are powerful normative reasons for pursuing more robust public engagement with science, from the need for public accountability for the values embedded in science to the need for the public to understand scientific practice and results. There has also been a plethora of experimentation with public engagement or "citizen science" approaches over the past few decades. (See, e.g., Douglas 2005, Brown 2006, Cavalier & Kennedy 2016, Ottinger 2017) However, there seems to be a growing skepticism that such public engagement attempts can deliver much of value. The ability of public engagement to shape science policy, to help inform the public about science, or to hold scientists accountable for value judgments has been criticized and now seems further out of reach than ever. (See e.g. Stilgoe, Lock, & Wilsdon 2014, and the papers in that special issue on public engagement.) We will argue in this paper that is in part because of an under-theorization of the public in such practices. The experimentation that has proceeded has not had a helpful ontology for how the public is constituted in these practices, and as a result, has not had a proper grasp on the normative aspects each sense of public brings with it. There are at least four different conceptions of the public that are utilized in public engagement activities (where engagement requires bi-directional interaction, and thus does not include linear communication efforts). They are 1) the public as volunteers (whomever is willing to participate); 2) the public as a representative sample (a carefully selected sampling, often following demographic parameters); 3) the public as stakeholders (who represent various interests); and 4) the public as a community. We will show how these conceptions of the public are distinct, and further show how the different conceptions bring different strengths and weaknesses to public engagement practices, including epistemic, ethical, political, and pragmatic strengths and weaknesses. By attending more carefully to the normative aspects of different conceptualizations of the public, public engagement with science practices can be better tuned to achieve what they can plausibly achieve. They are crucial components to the generation and utilization of expertise in democratic societies, but they must be deployed with attention to the way in which the public is conceived within each practice.

Brown, M. B. 2006. Survey Article: Citizen Panels and the Concept of Representation. *Journal of Political Philosophy,* 14(2).
Cavalier, D., & Kennedy, E. B. (Eds.). 2016. *The Rightful Place of Science: Citizen Science*. Consortium for Science, Policy & Outcomes.
Douglas, H. 2005. Inserting the Public into Science. In *Democratization of Expertise?* (pp. 153-169). Springer, Dordrecht.
Ottinger, G. 2017. Reconstructing or Reproducing? Scientific Authority and Models of Change in two Traditions of Citizen Science. *The Routledge Handbook of the Political Economy of Science*, 351- 363.
Stilgoe, J., Lock, S. J., & Wilsdon, J. 2014. Why Should we Promote Public Engagement with Science?. *Public Understanding of Science*, 23(1).

28. Working with Fossils: Stratigraphic Frameworks, Correlation and the Problem of Nomic Measurement
Max Dresow

This paper gives an unconventional answer to a familiar question. The question is, how is a science of geohistory possible at all? The question arises because a science of geohistory studies things that no longer exist, or at any rate, are no longer present in the same way that stars and starfish are: past states of the earth and its biological inhabitants. How then is it possible to conduct systematic inquiry into these (non-)entities?

There is a conventional answer to this question, and it is the principle of actualism. This is the 'principle' that the present is the key to the past—in other words, that processes active in the present supply the best guide for interpreting traces of past events and conditions. According to actualism, the way we learn about the past is by studying how causes in the present ('actual causes') produce their effects. It follows that a science of geohistory is possible insofar as scientists can (1) learn about how contemporary processes produce their effects, and (2) use this knowledge to make inferences about past events and conditions.

The conventional answer is a perfectly fine answer; yet it is not the only answer, and if we are interested in the organization and practice of specifically geohistorical science, it may not be the most interesting. Consider that the conventional answer says that a science of geohistory is possible because warrants permit inferential access to the deep past. But all inductive inference relies on warrants to move from evidence to conclusions. The conventional answer thus fails to illuminate anything distinctive about a science of geohistory.

I offer a different answer--one that begins with the observation that, in geohistorical research, placing events in their proper time position and perspective is central to a great deal of investigative practice. And here one epistemic resource is indispensable: a stratigraphic framework. Also called a geological time scale, a stratigraphic framework is an invaluable tool for geoscientists investigating virtually any aspect of Earth's development, anywhere in the world. Kant might call it a 'condition on the possibility of reconstructing geohistory'—at any rate, it is highly important.

But stratigraphic frameworks are not simple conventions; they are compilations of measurements, and as such, they need to be assembled using empirical data. This paper is about how this is done. Specifically, it is about the kinds of work involved in articulating an early draft of the geological time scale, as well as an important problem this project faced—the absence of a theoretical justification, at least at first, for the practice of using fossils to measure time. This problem is an instance of what Chang calls the "problem of nomic measurement": geologists want to measure one variable (time), but only have access to a second variable (faunal composition) whose relationship to the first variable is unknown. After giving a detailed characterization of the problem, the paper considers how it was navigated in ongoing practice: by an iterative process guided by judgments of heuristic utility.

29. Mental Suffering and Treatable Conditions
Andrew Evans

Some people have intense feelings of worthlessness and cannot get out of bed in the morning. Others experience hallucinations that prevent them from competing daily activities. Still others, despite their best efforts, cannot stop using heroin. We usually use the term "mental disorder" to refer to these, and other, conditions. In fact, psychiatrists, psychologists, social workers, counselors, and philosophers all use the term "mental disorder", and despite many attempts at articulating what the term means, there is no universal agreement. Further, many of the concepts of mental disorder that have been put forward by philosophers of mental health do not conform with clinical practice. So, what would a concept look like that better conforms with clinical practice? What would a concept look like that captures what the mental health community actually treats? In this paper I argue that in order to give a concept that conforms to clinical practice, the term "mental disorder" should be dropped all together. Using the work of Derek Bolton (2008), I argue that a concept of "treatable conditions", centered on mental suffering and the need to treat, is more in line with clinical practice.

I begin in section (2) with an overview of some of the concepts of mental disorder that have been put forward and identify threads that run through the accounts. Notably, harm and dysfunction seem to be the core of what we mean when we say "mental disorder". But there are conditions that are appropriate for treatment that do not involve dysfunction. So, in section (3) I argue that none of these concepts of mental disorder are a good match with clinical practice. We should therefore drop the term "mental disorder". However, we still need some conceptualization of the conditions that the mental health community treats. So, I develop a concept of "treatable conditions": conditions associated with mental suffering that are appropriate for the intervention of the mental health community. In section (4) I argue that this concept is a good match with clinical practice since a core feature of the mental health community is response to suffering. In section (5) I discuss some of the factors that go into determining whether treatment is appropriate, what sort of treatment is appropriate, and who gets to decide. I argue that these decisions about treatment appropriateness are made by a collaboration between relevant stakeholders, including patient and clinician. In section (6) I respond to potential objections and conclude in section (7) by comparing my account to Bolton's account (2008).

30. The Communicative Responsibilities of Scientists
Paul Franco

In 2015, Bjorn Stevens took what *Scientific American* called "the unusual step, for a climate scientist, of issuing a press release to correct…misconceptions" in popular media discussions about a study accepted for publication in the *Journal of Climate*. Part of these misconceptions could be traced to a blog post from the Cato Institute, which claimed Stevens's paper "finds the magnitude of the cooling effect from anthropogenic aerosol emissions during the late 19th and 20th century was less than currently believed, which eliminates the support for the high-end negative estimates (such as those included in the latest assessment of the U.N.'s Intergovernmental Panel on Climate Change, IPCC)."

In his press release, Stevens says that while his "paper presents a number of arguments as to why the radiative forcing from aerosols is neither as negative, nor as uncertain, as has previously been thought," his "new estimates of aerosol radiative forcing are within the range of previous estimates, e.g., as provided in Chapter 7 of the IPCC fifth assessment report…albeit on the lower end of that range in terms of the estimated magnitude of the forcing." Stevens goes on to point out his study did not support some of the inferences being made in certain corners of the media about the insensitivity of the Earth's surface temperatures to atmospheric CO_2. He cautions that "one should resist concluding too much, too soon, from a single study. In the long run I certainly hope that my findings will help constrain the climate's sensitivity to CO_2 but they do not, on their own, relieve society of the threat of dangerous warming arising from anthropogenic emissions of CO_2."

In this paper, I use Stevens's press release to think about how epistemic and nonepistemic considerations shape the communicative responsibilities of scientists. To do this, I consider recent literature in science and values calling for increased attention to the ways in which nonepistemic values might or might not have a role to play in scientific communication (John 2015, 2018, 2019; Franco 2017, 2019; Goodwin 2018). After situating my project relative to this work, I draw upon speech act theory to emphasize the importance of audience uptake in the successful performance of a speech act. In particular, I argue that securing successful uptake in one's audience involves preventing misunderstandings and addressing misinterpretations of both the force and meaning of a speech act (cf. Elliott 2017, ch.6, and Franco 2019). Further, insofar as these possible misunderstandings and misinterpretations can have wide-ranging nonepistemic consequences, it is part of scientists' communicative responsibilities, in an ethical sense, to take seriously the importance of securing successful uptake (cf. Douglas 2009, and Franco 2017).

Finally, time and space permitting, I'll draw upon Mary Kate McGowan's work in conversational kinematics to think about the ways in which actions like Stevens's press release, insofar as they help secure uptake in their audience, help shape both the content and norms surrounding public discussion of scientific results in positive ways.

31. Integrating Practices from Multiple Specialties: Some Problems to Address
Elihu Gerson

Research increasingly depends on integration of practices from multiple specialties. This suggests thinking about several distinctions that are seldom examined: between integration and intersection of specialties or research programs; between particular projects and programs; and between epistemic and organizational integration.
Intersection occurs when different programs overlap to form a new program that includes features of both. Studies of development, for example, are similar in many ways across the classes of vertebrates. But they are similar rather than integrated; their results do not

directly depend upon or influence one another. Instead, they simply share a larger intellectual context.

Integration occurs when multiple research programs develop complementary practices concerned with common problems, models, data handling techniques, or characterizations of phenomena. Successful integration requires different capacities from participating programs, which creates new or revised practices that don't necessarily appear in the parent programs. Hence, integration creates intersection, but not necessarily vice-versa. Integration is realized in particular joint projects that use resources from different programs to address common problems. Because projects involve particular studies, they require levels and kinds of effort which intersection doesn't. For example, if a concept in one program is inadequate or incompatible in the context of another, that is no barrier to continued research in either, although it may stimulate efforts to reconcile the difference. In a joint project, such difficulties must be reconciled if the work is to continue to a satisfactory end.

Because different programs supply complementary capacities, joint projects create common knowledge, i.e., knowledge which all participants can reliably assume that other participants have. This common knowledge and its associated practices that constitute epistemic integration. New common knowledge then supports new programs as they develop.

Programs or specialties and projects are thus very different. Projects require explicit efforts to identify and resolve technical difficulties that may not even be known in the separate programs joined in the project. Moreover, projects have additional characteristics which must be taken into account. In particular, they employ particular people and particular resources, and are housed and administered in particular organizations. This raises another problem with the notion of integration.

Many organizational arrangements (e.g., administrative support activities) vary with little respect to epistemic concerns. Varying organizational practices across different specialties, host institutions, or jurisdictions may impose additional burdens on projects, and thus distort or even block some research activities. Conversely, different epistemic requirements often require specialized organizational arrangements. Finding vertebrate fossils, e.g., requires different kinds of organizational arrangements than storing and curating them, and in some collections, the storage arrangements may be better suited to the needs of morphologists than, e.g, ecologist's, geneticist's, or systematist's.

Analyzing how integration creates and responds to intersections among programs thus requires considering how organizational practices shape and are shaped by epistemic considerations. In particular, at the moment of discovery, when practices are being formed and revised, the distinction between organizational and epistemic integration dissolves, and then re-forms as new practices become settled. These distinctions pose important problems for research on practices, their interactions, and their contexts.

32. Einstein's Practice of Principles. On Einstein's Distinction between Constructive and Principle Theories

Marco Giovanelli

Toward the end of 1919, in a two-column contribution for the *Times of London*, Einstein famously declared relativity theory to be a 'principle theory,' like thermodynamics, rather than a 'constructive theory,' like the kinetic theory of gases. This distinction has attracted considerable attention in both the historically- and the theoretically-oriented scholarship. As it turns out, its popularity in today's scholarship has somewhat hindered the appreciation of its core message. This paper hopes to show that to properly understand Einstein's 'theory of theories, ' one has to disentangle the two threads of its fabric, the context of justification, and the context of discovery. Einstein introduced not only a distinction between two types of existing theories but also between two strategies for finding new theories:

As an orientation, the paper organizes the available textual evidence in three successive phases.

- During the Swiss years (1905-1914), Einstein started to compare relativity theory to 'Thermodynamics before Boltzmann' as a negative defensive argument. In his 1905 relativity paper, Einstein had derived the dependence of the 'electron" s mass on velocity by adapting Newton's dynamics of charged point particles to the relativity principle. However, according to many (, this derivation was not admissible without making some assumptions about the structure of the electron. Einstein replied that the relativity is similar to thermodynamics. It does not directly try to construct models of specific physical systems but provides empirically motivated and mathematically formulated criteria for the acceptability of such theories.
- During the Berlin years (1914-1933), Einstein progressively transformed this line of defense into a positive heuristics that he applied in his gravitational research. Thermodynamics shows that instead of directly searching for new theories, it is often more effective to search for conditions that constraint the number of possible theories. Thus, Einstein introduced a sort of discovery algorithm: (a) search for empirical facts that can be translated into mathematically formulated principles; (b) check if the existing laws of nature are compatible with those principles (c) if not, modify them so that they do; (e) compare the modified laws with experience.
- During the Princeton years (1933-1955), Einstein presented on a few occasions the historical path to special relativity as an application of such heuristics. He suggested that he had consciously decided to develop relativity following the example of thermodynamics because of his skepticism towards the exact validity of Maxwell equations. The Lorentz transformations are not the 'byproduct' Maxwell equations or of any particular dynamical law; they are empirically grounded 'constraints' that all laws of nature have to satisfy.

The paper concludes that the characterization of relativity theory as a 'theory of principles,' cannot be disentangle from what might be called a 'practice of principles' (Seth, 2010). Principle theories are (a) a class of existing theories: principle theories do not entail single physical laws, put constraints on them (b) a strategy for finding new theories: instead of searching directly for the laws of nature, first search for constraints the limit the number of possible candidates.

33. What makes Explanations Good in Engineering Science?

Kristian González Barman

In this contribution I address the question ‘what makes for good explanations in engineering science?’

This research is motivated by two observations. First, while there has been much discussion in the last decade about the goodness of mechanistic explanations (Craver 2007, 2014; Kaplan and Craver 2011, Levy and Bechtel 2013), most treatments focus only on virtues or norms by which to evaluate mechanistic representations, such as precision, abstraction, generality, completeness, and accuracy, yet do not spell out what explanatory insights are gained when mechanistic representations meet such norms.

Second, these treatments by and large ignored the question of what makes for good mechanistic explanations in the engineering sciences. The results of this research offer new insights into both these issues. The key claim that I will argue for, and illustrate with empirical examples drawn from engineering science, is that the above-mentioned virtues or norms (plus some others) are adequate criteria by which to evaluate mechanistic representations in engineering science, if explainers are better able to track relevant counterfactual dependency relations when mechanistic representations meet these virtues/norms. Which virtues/norms conduce to the tracking of relevant dependencies, and which trade-off, depends on the specifics of the request for explanation, as the engineering cases clearly show. I will discuss a case from civil engineering on the collapse of a convention centre and one from mechanical engineering on broken wheel-shafts.

More specifically, in the cases I focus on, causal interactions in engineered mechanisms can be represented with structural equations, where some variables (of the structural equations) can be filled-in by symmetrical equations. For example, we might have that a wheel-shaft will break if a perpendicular force applied to its centre exceeds a certain amount (which can in turn be calculated using certain parameters -torque, mass, etc.- via symmetrical equations).

Good mechanistic representations enable the tracking of counterfactual dependencies through the choice, and the value-assignment (together with the calculation and further interpretation) of these variables. In the previous example, we could ask whether or not the shaft would break if it had been made with a material with a different tensile strength, or what would happen if the radius had been twice the original size. In order for an explanation to enable tracking the relevant counterfactual dependencies, the variables need to be appropriately chosen, and this is done by aiming at the maximization of certain virtues/norms (e.g., we can choose more fine-grained parameters in our equations to get a more precise result for the maximum resistance force).

This link between (the maximization of certain) virtues/norms and the tracking of relevant dependencies is relative to specific explanatory requests, which in turn is relative to specific engineering design aims. I show that robustness is a key norm for mechanistic representations in mechanical engineering, while precision is a key norm in civil engineering.

Craver, C. F. 2007. *Explaining the Brain: Mechanisms and the Mosaic Unity of Neuroscience*. Oxford University Press.

Craver, C. F. 2014. The Ontic account of Scientific Explanation. In *Explanation in the Special Sciences* (pp. 27-52). Springer, Dordrecht.
Kaplan, D. M., & Craver, C. F. 2011. The Explanatory Force of Dynamical and Mathematical Models in Neuroscience: A Mechanistic Perspective. *Philosophy of Science*, 78(4).
Levy, A., & Bechtel, W. 2013. Abstraction and the Organization of Mechanisms. *Philosophy of Science*, 80(2).

34. Understanding the Role of Representational Dissimilarity in Model Success
Deborah Haar

In *Simulation and Similarity* (2013), Michael Weisberg further develops the similarity view of modeling, in which the relationship between a model and its target is one of similarity in relevant respects and degrees. One commitment of this view is that model-target similarity underwrites the success of a model. Weisberg proposes a formal similarity measure, using the model's and target's attributes (states) and mechanisms (transition rules) as the points of comparison. While I do not dispute that model and target similarity at the constituent level is important to modeling success, there are cases from computer simulation literature which illustrate that component similarity may not always underwrite success; cases where there is representational dissimilarity at the level of attributes and mechanisms, but modeling success nonetheless. These examples cannot be accounted for by appealing to the kind of modeling taking place, or by appealing to abstraction or idealizations. This suggests that model success cannot be solely attributed to some form of representational similarity at the level of attributes and mechanisms.

For an example of how attribute and mechanism similarity may not underwrite model success, we can look to Johannes Lenhard's 2007 paper which presents a case study of the development of the first simulation of atmospheric dynamics. The simulation was deemed a general success because it reproduced the patterns found in the real-world system. However, after a month the simulation of stable wind patterns 'exploded' into chaotic behavior. The problem had computational origins (a build-up of truncation errors) and the solution that was eventually accepted, proposed by Akio Arakawa, had implications that were incongruent to the known properties of the physical system. The result of the fix is that the model behaves like the target, but the mechanisms that constituted the fix were not similar to the target.

This example shows that an analysis of model and target parts for similarity (or even 'feature sharing', see (Khosrowi 2018)) may not serve as the basis for an estimation of or an explanation for a model's success, at least when it comes to computational modeling. In addition to any similarity analysis done at the level of the model's and target's parts, we need to consider the behavior of the model when it is run. Functional similarity, or similarity of the behavior of the model to its target, is not guaranteed by or is the same as component similarity. Functional similarity allows for us to account for computational elements of the model that would not be included in a formal analysis of the relative similarity of states and mechanisms of a model and its target. Simulation modelers are aware not only of the similarities between states and mechanisms of their model and the system, but of the mathematical and computational techniques utilized in the model and any errors that arise when the simulation is run. These computational aspects are, of course, not found in the target system, but they play a key role in the ability of computational models to successful represent physical systems.

35. Leaving the Box Closed: An Interventionist Approach to Machine Learning for Drug Design
Hyejeong Han

The recent advancement of Artificial Intelligence (AI) and Machine Learning (ML) techniques has seen them take a central role in the practice of science. While ML shows high predictive accuracy, its low explanatory power has been pointed out as a major problem. Why is explanation needed and how is explanation possible in the use of ML for the practice of science? In this talk, I will present an interventionist approach to ML for the practice of science by investigating a case of Quantitative Structure-Activity Relationship (QSAR) modelling for drug design.

Drug design, or rational drug design, refers to a drug discovery strategy that finds a new drug efficiently based on a systematic application of chemical and biological knowledge. A statistical description of the relationship between the biological activity of a molecule and its chemical structure, so called QSAR modelling, is one of the central methods for drug design. The first modern QSAR model was invented in the early 1960s in search of molecular mechanism of plant growth control, and many variations of QSAR models have been typically used for explanatory purposes. However, a new trend which pursues prediction over explanation boomed in the QSAR community since the 1990s, and is characterised by the reliance on ML techniques. Centring around the trade-off between predictive accuracy and explanatory power of ML, there have been tensions between the “classical” QSAR modelling group and ML-assisted QSAR modelling group (Fujita & Winkler, 2016). While the ML-assisted QSAR modelling group has prioritized predictive accuracy, many efforts to enhance the explanatory power of ML-assisted QSAR models have continued as well.

In this talk, I will critically review the efforts for explanatory ML-assisted QSAR model and explicate the proper meaning of explanation in the context of drug design. In the use of a scientific model, the meaning of explanation can be divided into first, the explanation about the working mechanism of the model, and second, model-based explanation about the target phenomenon. Based on the distinction of the two different meanings of explanation, I will defend both the necessity and possibility of the second type of explanation in ML-assisted QSAR model. More specifically, I will present how ML-assisted QSAR model can and should generate explanation for drug design without opening the “Black Box”, i.e., leaving it opaque how the complex ML-assisted model works. The defence for explanatory ML-assisted QSAR model will be based upon the interventionist approach suggested by James Woodward (2003).

This talk contributes to study the modelling practice of science from a different perspective. The emerging usage of ML techniques adds a novel dimension to the philosophy of scientific models, since modelling is practised at least partially by the machine based upon an iterative process without being explicitly programmed by a human. Therefore, human centred approaches to scientific model and explanation may not be sufficient to understand the emerging practice of scientific modelling. In that sense, this talk contributes to the philosophy of scientific models from a novel perspective by examining explanation in ML-assisted practice of science.

Woodward, J. 2003. *Making Things Happen: A Theory of Causal Explanation*. New York: Oxford University Press.
Fujita, T., & Winkler, D. 2016. Understanding the Roles of the "Two QSARs". *Journal of Chemical Information and Modeling*, 56(2).

36. Why Simpler Computer Simulation Models can be Epistemically Better for Informing Decisions

Casey Helgeson, Vivek Srikrishnan, Klaus Keller and Nancy Tuana

Computer simulation models are now essential tools in many scientific fields, and a rapidly-expanding philosophical literature examines a host of accompanying methodological and epistemological questions about their roles and uses. Climate science is one such field, and questions about the interpretation and reliability of the simulation models used to understand, attribute, and predict climate change have received considerable philosophical attention.

One conspicuous feature of scientific discourse about the simulation models used in climate science, and in environmental modeling more broadly, is the attention given to where a model lies on a spectrum from simple to complex. While this attention to model complexity has informed some of the philosophical discourse on simulation modeling, its relevance for the literature on simplicity in science remains largely unexplored.

This literature on simplicity addresses whether and why simpler theories (or hypotheses, models, etc.) might be—other things being equal—better than complex ones. Different ways of unpacking ``simpler'' and ``better'' yield a diversity of specific theses, with correspondingly different justifications. A number of distinctly modern variants rest on mathematical theorems tying well-defined notions of simplicity to benefits such as predictive accuracy, reliability, and efficient inquiry.

Here we discuss a notion of simplicity drawn from scientific discourse on environmental simulation modeling and expound its importance in the context of climate risk management. The new idea that we bring to the simplicity literature is that simplicity benefits the assessment of uncertainty in the model's predictions. The short explanation for this is that quantifying uncertainty in the predictions of computer simulation models requires running the model many times over using different inputs, and simpler models enable this because they demand less computer processor time.

While complexity obstructs uncertainty quantification, complex models may behave more like the real-world system, especially when pushed into Anthropocene conditions. So there is a trade-off between a model's capacity to realistically represent the system and its capacity to tell us how confident it is in its predictions. Both are desirable from a purely scientific or epistemic perspective as well as for their contributions to the model's utility in climate risk management. Whether simpler is better in any given case depends on details that go beyond the scope of this paper, but the critical importance of uncertainty assessment for addressing climate risks is why simpler models can be epistemically better for informing decisions.

In this paper/presentation, we first introduce the relevant notion of simplicity and a way to measure it. Using detailed examples of current statistical practice applied to simulation models of the Antarctic Ice Sheet, we explain and illustrate the link from simplicity to uncertainty quantification. We argue that through this link, simplicity becomes epistemically relevant to model choice and model development. We discuss the resulting trade-off, highlighting the roles of non-epistemic values and high-impact, low-probability outcomes in mediating the importance of uncertainty assessment for climate risk management.

37. Where the Wild Things are Classified: Consequences of Classifying Wildness and Nonhuman Animals Caught in the Crosshairs
Denise Hossom

A common understanding of what makes something “wild” existed long before scientific criteria to quantify and qualify wildness, or the adoption of wildness as the focus of conservation efforts. On the common understanding, we don’t need biological sciences to identify “wild” things, especially nonhuman wild animals. Yet our scientific classification practices for wild nonhuman animals are becoming increasingly complex due to the breadth and range of human – nonhuman animal interactions and relations. From zoology to invasion biology, classifying nonhuman animals as “wild” is getting complicated.

In this paper I report on classification practices for nonhuman animals in terms of five interrelated distinctions; 1) Wild – Domestic, 2) Feral – Tame, 3) Free – Captive, 4) Indigenous – Invasive, 5) Vermin/Varmint/Pest – Livestock/Pet. I draw on case studies from my fieldwork in progress with scientists engaged in conservation of snow leopards in Nepal and at San Francisco Zoo, and scientists supporting population control measures of wild horses – mustangs in the U.S. and brumbies in Australia. One aim of my account is to provide an analytical schema, implementable in practice, for particular cases of conflicting scientific and social aims.

Conservation projects, like much of science, are practiced within a social context and therefore intersects with both scientific and nonscientific values. I am interested in conflicts where both scientific and social goals are at stake in how a conflict is resolved. The analytical schema I propose is a tool for scientists, aimed at facilitating the identification of individuals and groups of humans, nonhuman animals, and even material or environmental stakeholders that are engaged in the conflict.

Each case study is examined schematically as a “zone of conflict”, where the intersections of animals classified as wild and non-wild create tensions between stakeholders.
An example of a “zone of conflict” that illustrates the complexity of classification practices and its consequences is horses, inclusive of the common “domestic” horse (E. f. caballus), and the rarer “wild” type of Przewalskii horse (E. f. przewalskii). Contentions over how to classify each of these as species or subspecies is ongoing, and both varieties of horse have found themselves on every side of the five distinctions, both historically and contemporarily. The classification of horses is often a fluid and shifting practice and can be highly context dependent. The case of classification practices for American Mustangs and Australian Brumbies demonstrates ongoing conflicts between scientists, government agencies, and the broader public on how to classify and manage their populations, and difficult policy decisions and implications. The schema I propose develops through three stages: Stage one outlines the background context for each stakeholder with their corresponding interests, values, aims, and practices. Stage two frames how each stakeholders background

context either conflicts or aligns with other stakeholders. Stage three is the development of resolution practices.

38. Making it Manifest: The Methodological Significance of Good Variable Choices
Josh Hunt

Philosophers of science have recently broached an important methodological question: why are some choices of variables better than others? Answering this question is essential for understanding how scientists and engineers make progress by reformulating their theories. Although we could "in principle" solve many problems with a poor choice of variables, in practice we require good variable choices to make problems tractable and computationally feasible. The need for good variables arises whenever practitioners apply mathematics to formulate their models or theories. Hence, the question of good variable choice is pressing not only for philosophers of the natural and social sciences, but also for philosophers of engineering.

From the outset, there is a worry that any satisfying answer to the question of variable choice will embroil us in metaphysical speculation. Metaphysicians such as David Lewis and Ted Sider have argued that certain formulations of scientific theories are more fundamental than others because they use perfectly natural or "joint-carving" properties. Indeed, when scientists discuss good variables, they often extol their ability to make certain properties manifest. This is especially common with systems that display a high degree of symmetry, prompting variable choices that make the relevant symmetries manifest. At first glance, this notion of making properties manifest seems to be grist for the mill of the metaphysical joint-carver while being anathema for more empirically-minded philosophers. Fortunately, we can provide a compelling methodological account of the significance of making properties manifest. Using a variety of case studies from physics, chemistry, and engineering, I will argue that there are two key components to this methodological account.

First, there is an epistemic component: by working in variables that make a property manifest, scientists minimize what they need to know to determine if a given expression has the property of interest. For instance, in particle physics, practitioners typically choose variables so that the Lagrangian becomes manifestly local and Lorentz invariant. One can then tell by inspection whether or not a given expression is local or Lorentz invariant. Furthermore, good variable choices typically preserve the property of interest across standard calculational manipulations. This ensures that the calculated expression possesses the relevant property without needing to perform further checks.

Secondly, there is a pragmatic component: good variable choices simplify the calculations involved, often dramatically. Bad variables typically lead to convoluted expressions that nevertheless cancel through further manipulation. Good variable choices allow us to avoid working with these inconvenient or redundant expressions in the first place. The recent introduction of spinor–helicity variables in quantum field theory provides one dramatic example. This variable choice replaces hundreds of pages of calculations with simple three-page inductive arguments. By making problem-solving more convenient, good variables enable scientists to figure out how to solve problems in practice rather than merely in principle.

Further examples illustrating my account include elementary coordinate transformations common throughout science (such as using spherical coordinates for systems with spherical symmetry), Laplace and Fourier transform techniques in engineering, symmetry-adapted basis choices in quantum chemistry, and diagrammatic methods in physics and chemistry. Importantly, these examples illustrate how the advantages of good variable choices nevertheless come along with methodological disadvantages for other goals that practitioners have. These methodological trade-offs problematize the joint-carver's dream of a canonical language for describing science. Instead, scientific practice furnishes a wellspring of alternative variable choices that accomplish our goals more or less well depending on the properties they make manifest and the properties they obscure.

39. Canons of Algorithmic Inference: Feminist Theoretical Virtues in Machine Learning
Gabbrielle Johnson

As inductive decision-making procedures, the inferences made by machine learning programs are subject to underdetermination by evidence and bear inductive risk. Previous attempts to address these challenges have been guided by the presumption that machine learning processes can and should be formally objective. In doing so, the influence of values has been restricted to data and decision outcomes, thereby ignoring internal value-laden choice points. In this paper, I argue that these efforts rest on a mistake: the resources required to respond to these challenges render machine learning processes essentially value-laden, and thus sanction ethical and socio-historical interventions throughout their production, use, and evaluation. I demonstrate these points in the case of recidivism algorithms, arguing that the adoption of feminist theoretical virtues supports the use of false positive equality as a measure of fairness in order to stymie the ongoing harm to the black community within the criminal justice system.

40. On Politicization of Nutrition Guidelines
Saana Jukola

Government-issued nutrition guidelines (e.g., The Dietary Guidelines for Americans) guide health policy making, offer guidance to nutrition professionals and give information on what diets are considered to promote health and prevent chronic diseases at the population-level. Issuing these recommendations is a complex task involving not only amalgamating evidence from different studies but also taking into consideration diverse political, ethical, and cultural concerns. Many critics have questioned both the claim that nutrition guidelines are science-based and the political justification for publishing guidelines that influence national food programs and public meals (e.g., Kuttner 2014). According to these complaints, government-issued guidelines are not trustworthy.

This paper focuses on entanglement of science and political concerns in nutrition guideline development. How to guarantee the public trustworthiness of the guidelines when experts are accused of politicizing dietary recommendations? In particular, it addresses the controversy concerning the suggestions that dietary guidelines should recommend plant-based diets – not only on the basis of reducing the risk of chronic diseases and improving population health but also in order to promote sustainability (Freidberg 2016). According to the critics, including sustainability in the guidelines does not relate to the goal of promoting overall population health and, thus, is a sign of politicization of the guidelines (e.g., Kuttner 2014).

The first objective of the paper is explicate how the participants in the debates perceive the goals of population-level nutrition guidelines, especially the goal to promote population health. I will analyze the conceptions of (population) health that are used in the controversy. I will show that the critics of including sustainability claims use a very narrow concept of health and excessively strict understanding of determinants of health. The second aim is to address the value-ladenness of the guidelines and examine their legitimacy in contexts where the target population holds diverse and often conflicting ethical and political values: Guidelines are health policy measures that impose a considerable lower impact on individuals' freedom of choice than measures such as restrictions on manufacturing of certain foods (Resnik 2015). Yet, the guidelines have a practical implications for the lives of the laypeople. This gives rise to potential conflicts between the values implicit in the guidelines and the values of the target population, and, thus, threatens the trustworthiness of the guidelines. Examining how guidelines could be developed in a way that befits a wide variety of values in the target population is a question that has to be addressed. I will argue that the use of citizens panels in the process could help to reduce value conflicts.

Freidberg, S. 2016. Wicked Nutrition: The Controversial Greening of Official Dietary Guidance. *Gastronomica: The Journal of Critical Food Studies*, 16(2), 69-80.

Kuttner, H. 2014. *How to Sustain Sound Dietary Guidelines for Americans: Mission Creep within the Federal Dietary Guidelines Advisory Committee threatens Americans' Health and Well-being.* Hudson Institute, Washington

Resnik, D. B. 2015. Food and Beverage Policies and Public Health Ethics. *Health Care Analysis*, 23(2), 122-133.

41. Reconsidering the Werner-Jørgensen Controversy: A Crucial Experiment?
Kayoung Kim

The controversy between Alfred Werner(1866-1919) and Sophus Mads Jørgensen(1837-1914) on the structure of complex inorganic compounds is a well known case to chemists, owing to the extensive work of George B. Kauffman. In inorganic chemistry textbooks, the controversy has been introduced as a clear-cut case of Werner's triumph over Jørgensen supported by sufficient experimental evidence. However, considering that the notion of "crucial experiment" has been regarded suspicious to philosophers of science after Duhem, I suggest that this case needs a critical examination. Especially, a historian of science, Helge Kragh (1997) argued that the controversy was not epistemically resolved but rather abandoned, considering the historical context in which Jørgensen is said to have been admitted his defeat. If this controversy has room for reconsideration, can we say that, pace Kauffman, Jørgensen could have maintained his theory without loss of rationality even after Werner's 1907 experiment of violeo salt synthesis? In the paper, I will first reconstruct the case to exemplify a crucial experiment following Kauffman's historiography and then reconsider the case by focusing on some issues such as: the possibility of error in the experiment; the context of Werner's reliance of "isomer counting" technique; Werner's assumption of the valence of nitrogen, etc. Following Joseph Pitt (2001), I suggest that whether the case exemplifies a crucial experiment or not depends on selecting a relevant explanatory framework for the point that the historian wants to make.

42. Modeling Trade-Offs in Predictive Toxicology QSARs
Frederick Klaessig

In computational models, a claim is expressed in a language of certainty, a mathematical expression and its associated statistical methods. The current practice in quantitative structure activity relationship models (QSAR) is to validate fitting constants using half of the dataset in order to 'predict' the toxicity of the other half. Regulatory concerns that the resulting models are simply correlations, has led to the OECD criterion that there be a "mechanistic interpretation, if possible."

This was not the case in the historical development of QSARs which utilized thermodynamic and kinetic reaction rate concepts to frame the analysis, i.e. set constraints on the mathematical expressions. A multi-step reaction pathway leading to toxicant-enzyme receptor interactions yielded a three parameter model. In the pharmaceutical industry, the pursuit of greater precision and realism in QSAR models has led to >2,000 'descriptors' that effectively narrow the range of mechanistic explanations. The classical, extra-thermodynamic style was displaced by another style in a manner analogous to Levins's trade-offs in generality, realism and precision.

There are differences. The first use of QSARs was for pesticidal agents applied to the whole organism, and drug discovery enzyme target are developed separately from drug design. Nevertheless, QSARs-as-chemical-groups guide decisions, especially in the newer field of predictive toxicology, regarding the selection of in vitro assays or of testing controls or for read-across exercises to fill data gaps. QSARs also guide the development of physiologically based (PBTK) models that relate the applied dose (or exposure) to the the in vivo organ-level dose.

Finding the intersection of test assay results, adverse outcome pathways and organ-level dose is a major goal of international initiatives in predictive toxicology (Tox21), safer-by-design, and responsible research & innovation (RRI). Should the level of generality, realism and precision for toxicology differ from that for therapy, then basing toxicity modeling on therapeutic QSAR and PBTK models will be unsuccessful and the current validation criteria inappropriate. Further, if Levins's 'strategy for model building in population biology' also applies to QSAR modeling, then the analysis of natural selection for populations with commensurate demographic & evolutionary times for reproduction should have parallels to organs experiencing adverse outcomes threatening survivability.

Claim, style, and strategy are terms not often used in describing research programs, repertoires, paradigms, frameworks and scaffolds. Elements of successful claims and strategies anticipate their later use in theory; remnants of unsuccessful ones are memorialized in the introductions to past journal articles. For this discussion, the predictive toxicology research program has computational modeling as a strategy that may lead to a multiplicity of QSAR styles.

QSAR and PBTK computational models used in setting worker exposure limits for silver nanoparticles will be discussed in the context Levins's trade offs. The immunological system & survival displace population & reproduction relative to natural selection.

43. Computational and Conceptual Analyses (applied) to Stress Interdisciplinary Research Trends: An Articulation of History and Philosophy for Contemporary Science

Jan Pieter Konsman

The notion of stress has been used in several disciplines, including material science, physiology, neuroscience and psychology, and may thus have enabled cross-fostering between disciplines. Whether stress should be considered as an interfield theory serving as a bridge between fields of research in physiology, neuroscience and psychology or as a field of research itself with “a central problem, a domain of items taken to be facts related to that problem, general explanatory facts and goals providing expectations as to how the problem is to be solved, [and] techniques and methods” (Darden and Maull, *Phil. Sci*., 1977, p. 144) remains a question. However, when one accepts that “[t]ypically, several fields…comprise a discipline” and that “[s]ocially, these units are characterized by such features as having academic departments … national professional organizations, [and] journals” (Bechtel, *Biol. Philos*., 1993, p. 281), then many interdisciplinary research fields can be considered to have become disciplines, including psychoneuroendocrinology, psychoneuroimmunology and to some extent stress research. In addition, stress can be thought of as a bridging principle in psychoneuroendocrinology and psychoneuroimmunology.

More than 25 years ago, it was noted that stress was initially “investigated in different disciplines without much contact” before “research on stress [became] one of the areas in which much interdisciplinary work is done” (Van der Steen, *Biol. Philos.*, 1993, p. 263). In this context, efforts were undertaken to “work with one integrative stress concept” in which "stress" came to be defined as a particular relation between stimuli, internal states, and responses of organisms” (Van der Steen, *Biol. Philos*., 1993, p. 264). It is this kind of stress concepts that van der Steen criticized at the time arguing that “to make empirical research sensible, independent definitions of stimuli and responses are needed” (Van der Steen, *Biol. Philos.*, 1993, p. 264).

Given that today several journals and institutions still contain the word stress, one may wonder as to what kind(s) of concept(s) stress refers to, if the distinction between physiological and psychological stress is still operational and whether the notion of stress has been extended and linked to other notions. Here, these questions were addressed by investigating how concepts of stress have been used in interdisciplinary research by identifying publications on stress and (psycho)neuroendocrinology or (psycho)neuroimmunology and determining trajectories of conceptual change over the past four decades. Bbibliometric approaches were used to computationally establish research fronts, “as the state of the art of a line of research”, and their intellectual bases in the form of “citation and co-citation footprint[s]” (Chen, *J. Am. Soc. Inform. Sci. Technol*., 2006). On these bodies of literature, conceptual analysis was applied to determine to what stress refers and the conditions under which an entity is classified as stress. It is concluded that, while over the past 25 years stress refers more and more to the cellular level, psychological constructs such as coping and resilience have also been implemented in behavioral neuroscience. So, without allowing full integration, stress does seem to continue to enable cross-fostering between disciplines.

44. A New Technique, A New Knowledge, Why Not a New Know-how?
Chia-Hua Lin

Mathematical constructs developed to advance knowledge in one discipline are sometimes applied to study a different subject in another discipline. Prominent examples include game theory (see Grune-Yanoff 2011, 2016) and formal language theory. Formal language theory is a study of mathematically defined languages, initially formulated in the 1950s by Chomsky (1956, 1959) to investigate syntactic regularities of natural languages. Today, formal language theory remains a branch of mathematics and a subfield in linguistics (Levelt 2008/2018). Yet, as much of its subsequent development was achieved in theoretical computer science (Greibach 1981), it enjoys a more dominant presence in computer science than in linguistics. The goal of this paper is to show how the cross-disciplinary applications of formal language theory may enrich—and resolve the tension in—the recent philosophical discussion on model transfer and knowledge modification.
Philosophers of science have been analyzing the cross-disciplinary use of mathematical constructs in terms of knowledge transfer (e.g., Humphreys 2004, 2018; Herfeld and Lisciandra 2019; Castellane and Paternotte 2019). While the term "transfer" suggests a conservative process in nature, various case studies suggest that modifying the body of knowledge—or the subject matter to which it applies—is indispensable to a satisfactory cross-disciplinary transfer (e.g., Herfeld and Lisciandra 2019). To resolve this tension between transfer and modification, one may expand the notion of transfer to accommodate creative model transfer (e.g., Houkes and Zwart 2019). Another option is to reconsider whether the term 'transfer' is adequate in characterizing the cross-disciplinary use of mathematical constructs (e.g., Grune-Yanoff 2011). Following a pragmatist view that takes models as epistemic tools (Boon and Knuuttila 2009), this paper presents a third approach. I argue that the oxymoron of modifying knowledge while transferring might be due to a conflation between different kinds of knowledge: knowledge about the mathematical content of a theory (or 'mathematical knowledge') and knowledge about the techniques ('know-how') of using the mathematical concent of the theory. Suppose that empirical knowledge and mathematical knowledge are two mutually exclusive sets of statements, noted as M={m1, …, mn} and E={e1, …, en}. Suppose also that a particular subset of mathematical knowledge (e.g., non-Euclidean geometry), when put in the right theoretical context, helps to establish knowledge about a certain empirical aspect of the world (e.g., Einstein's general relativity theory). A technique of using the mathematical content of a versatile theory thus refers to a particular way of converting a subset in M to a subset in E. To illustrate, I discuss a progression of examples from linguistics, theoretical computer science, and experimental psychology. These examples show that developing new know-how of using formal language theory does not always coincide with discovering new mathematical content. The tension between transfer and modification can thus be resolved by noting that generating knowledge of one kind and conserving knowledge of another kind can just be two sides of the same coin.

45. Null Hypothesis Statistical Testing and Psychology – a Case of Bad Scientific Practice
Matthew Lund

One tacit assumption of much SPSP work to date has been that contemporary scientific practice is epistemically significant, and that practice embodies some of science's epistemic norms. It is not common for the practice of some scientific field to be judged as epistemologically defective. However, if practice is a significant part of science and science can go wrong, practice could well be the culprit. This paper argues that psychology's practice is epistemically defective. While psychology uses many of the same statistical tools as other fields, it currently lacks the practical feedback loops to ensure that its tools and data are used responsibly. If practice is defined as "organized or regulated activities aimed at the achievement of certain goals", we can understand the practice of psychology as being centrally concerned with the goal of maximizing publication rates, eschewing thereby many traditional epistemic practices.

For much of the past decade, the field of psychology has been rocked by a replication crisis. (Pashler and Wagenmakers, 2012) Many published results that met the standards for significance, as defined by psychology's methods, have proved not to be reproducible. (Nuzzo, 2014) There have been two main theories of reform within psychology. The first blames most of the problematic features of the discipline on dishonest or questionable research practices (QRPs). For instance, there have been notable cases of individual fraud; also, there have been many cases of hiding data that does not lead to a certain p-value for a hypothesis and other instances of so-called p-hacking, i.e. manipulating the data set until it supports a given p-value. Another line of critique calls into question the legitimacy of Null Hypothesis Significance Statistical Testing (NHST) itself. Such critiques explore the dubious historical origins of NHST and show how the p-value is not – contrary to the opinions of many researchers – the touchstone to replicability and other notable theoretical virtues.

While there certainly are many valuable lessons to be learned from both of these lines of inquiry, they do not take into sufficient account the practical environment of psychology. All scientific fields suffer, in varying degrees, from fraud and questionable research practices. Moreover, many other fields use the statistical tools of NHST in much the same way as psychology, but such fields – with a few notable exceptions – have not been vitiated by abuses to the extent of a crisis.

Here is a list of practice defects endemic to contemporary psychology:

1) The majority of psychologists interpret central statistical variables incorrectly. For instance, the majority of psychologists surveyed committed to at least one of the following false beliefs: a. that a study having a p-value below 0.05 implies that the chance of replicability is 1-p, b. that a p-value below 0.05 proves the reality of an effect, or c. the probability of the alternative hypothesis (to the null hypothesis) being true is 1-p. (Gigerenzer, 2018)
2) The collection and evaluation of data is ordinarily done by the same (interested) parties. This is in contrast, for example, to NIH sponsored clinical trials, which require evaluation by an independent statistician as part of the research process.
3) Having a p-value below 0.05 is generally an essential condition for publication, and other aspects of the experimental situation are therefore not considered important.
4) As a field, psychology values the publication of research most highly. Psychology is not particularly unified, and new studies are largely independent of other work in the field. Thus, should a study reporting false results be published, it will be unlikely to conflict with extant studies, and hence will not be detected through consistency testing.

This paper advocates that psychology borrow some “best practices” from other fields to alleviate its crisis. For instance, p-values ought not to be the sole determinant of whether a study is publishable, statistical evaluation should be done by independent statisticians, psychologists should be better trained in statistical methods, full data sets of studies ought to be published, and preprints of studies should be encouraged.

Gigerenzer, Gerd. 2018. Statistical Rituals: The Replication Delusion and How We Got There. *Advances in Methods and Practices in Psychological Science*. 1 (2).
Nuzzo, Regina. 2014. Scientific Method: Statistical Errors. *Nature*. 506.
Pashler Harold, and Eric-Jan Wagenmakers. 2012. Editors' Introduction to the Special Section on Replicability in *Psychological Science: A Crisis of Confidence? Perspectives on Psychological Science : A Journal of the Association for Psychological Science*. 7 (6).

46. Epistemology of Science and Data-driven vs Hypothesis-driven Science
James Marcum

Contemporary data-driven or discovery science is often contrasted with traditional hypothesis-driven science. Advocates of the former champion it as the future of scientific practice and knowing because of the shear complexity that many natural phenomena exhibit. In other words, the latter is limited in terms of formulating hypotheses that can be subsequently tested in order to provide comprehensive theories to account for these phenomena sufficiently, if not completely. Unfortunately, as the history of science testifies, such theories are few, if not, non-existent. On the other hand, data-driven science, as its proponents insist, is not subject to this limitation. Rather, through analysis and modeling of big data obtained from quantifying (exhaustively) natural phenomena, an account of these phenomena is possible without resorting to either hypotheses or theories. However, advocates of hypothesis-driven science critique data-driven scientists as being unable to provide meaningful interpretation of their big data without an explanatory theoretical framework, which only hypothesis-driven science can provide. But, zealots of data-driven science, like Chris Anderson, who, in a 2008 issue of *Wired* Magazine, charged that big data signals the end of theories themselves in that the data can “speak” for themselves. Is there a way pass this stalemate?

In this paper, I propose a complementary model between data-driven and hypothesis-driven science. To that end, I use an epistemology of science to analyze both approaches to scientific practice and knowledge. Epistemology of science is concerned with the epistemic framework within which science is practiced and scientific knowledge is produced. Its elements include, for example, the nature and roles of hypotheses and theories, as well as epistemic assumptions, values, virtues, and norms, in practicing science and in producing scientific knowledge; the nature of scientific understanding and explanation; the structure of scientific-epistemic communities and the division of labor and agency within them; and the experimental and epistemic practices of scientists, along with the nature of scientific evidence itself; among other elements. From this analysis, rather than contrasting the elements of the epistemic framework for data-driven and hypothesis-driven science in binary oppositional terms, a complementary approach for integrating the elements of these approaches to scientific practice is taken, which provides a robust way for understanding it. For example, the epistemic assumptions of reductionism and holism or the epistemic values of simplicity and complexity complement each other as science analyzes the components of natural phenomena and then strives to reassemble those components in order to account theoretically for those phenomena. Briefly, the complementation consists of a yin-yang relationship between data-driven and hypothesis-driven science in which scientific practice can begin with hypotheses to be tested but then proceeds to the data, which shape hypotheses and, in turn, the data are subsequently shaped by the hypotheses. But, scientific practice can also begin with amassing large amounts of data, such as from omics technologies, which are then analyzed with statistical algorithms for epistemic patterns. In traditional formulary terms, empiricism (of data-drive science) and rationalism (of hypothesis-driven science) are two sides of the same epistemic coin.

Finally, the paper examines the epistemic issues that have emerged over the past several decades concerning the philosophical and epistemological issues surrounding data-driven science. One important issue is whether this science is changing the epistemic landscape from a traditional approach based on rationalism to one based on a new empiricism.

47. Can we use a Contradictions Right From The Start Methodology?
Maria Del Rosario Martinez Ordaz and Moises Macias Bustos

Here we tackle the question under which circumstances, if any, should physicists adopt a methodology that accepts a contradictions right from the start in order to achieve better understanding of the quantum phenomena?

Da Costa and de Ronde (2013) have argued in favor of developing an interpretation of superposition which, 'right from the start', takes contradictions to be a privileged element of the structured of Quantum Mechanics. According to them, a contradictions right from the start methodology can help physicists to take seriously the features which the theory seems to show, and with it, to explain out some of the alleged anomalies of the theory and its interpretations.

If they are in the right, there would be two important outcomes associated to their proposal: on the one hand, to adopt a methodology of this kind would help physicists to achieve better understanding of the quantum realm. On the other hand, philosophers would have helped physicists do develop a novel approach to their object of study, and this would reinforce the idea that philosophy has an important impact in the development of the sciences. The combination of these facts leaves us with the impression that the study of proposals such as the contradictions right from the start methodology deserve significant attention. Hence the importance of addressing this issue here.

In what follows, we describe in detail what a contradictions right from the start methodology should be and we provide a general guide for adopting such a methodological approach in quantum physics. In order to do so, we proceed as follows: First we introduce the contradictions right from the start methodology as it was presented by da Costa and de Ronde, then we extend it into a more cohesive proposal of what this methodology should be in order to be relevant for the physicists' practice. Second, we challenge the scope of this methodology by using it to tackle one of the most important problems of the GRW dynamical-collapse theory, namely, the 'problem of tails' (Albert and Loewer 1990, Wallace 2014). We reconstruct in terms of contradictions the problem of tails (in the form of both the problem of bare tails and the problem of structured tails). Third, we evaluate this methodology's usefulness for explanation of this problem and contend that while the approach promises to enhance our explanatory power it's not straightforward how it applies to these cases and whether it succeeds. Finally, we draw some remarks on the contexts in which physicists could (and should) adopt this methodology in order to achieve better understanding of the quantum phenomena.

48. Applying Unrigorous Mathematics
Colin McCullough-Benner

It is common, particularly in physics, for scientists to employ mathematical tools that do not meet the standards of rigor of pure mathematics. While in some cases this lack of mathematical rigor is clearly harmless, in others it is not clear that the mathematical objects that scientists seem to work with are even well-defined. In this paper, I argue that the most common philosophical account of applications of mathematics, the mapping account, can provide at best a misleading picture of the practices involved in applying unrigorous mathematics of the second kind. I then sketch a better way to represent these practices.

According to the mapping account, mathematical scientific representations represent their target systems as bearing a structural similarity to a structure picked out by the mathematics, with this similarity cashed out in terms of a structure-preserving mapping between the structure of the target system and the mathematical structure. Scientists' mathematically mediated inferences are then justified by the existence of such a mapping.

The first difficulty is arriving at a well-defined mathematical structure that correctly captures the accuracy conditions of the representation, since in these cases it is not clear that scientists are working with well-defined mathematical structures. But even once we have a structure that can do this work, the mapping account obscures the central features of the practice of applying unrigorous mathematics. In particular, applications of unrigorous mathematics are typically successful—when they are—thanks to inference strategies that confine problematic mathematical concepts to contexts in which they are well-behaved (Davey 2003). I illustrate this and the points that follow with two cases: Heaviside's use of his operational calculus to solve transmission-line problems and the use of path integrals in quantum field theory.

These inference strategies are not apt to be explained in terms of mappings for two reasons. First, in contexts in which scientists work with problematic mathematical concepts, so that it is not clear that they are reasoning about well-defined structures, an appropriate structure/mapping pair to represent this practice in terms of the mapping account is one that makes the inferences allowed by the inference strategy come out as truth-preserving. Since we have no grip on such a structure in these cases independently of the inference strategies scientists actually use, it is misleading to present such a mapping as underwriting the relevant inferences. Second, these inference strategies are frequently messy in ways that cause trouble for the mapping account. For instance, Heaviside appeals to the physical interpretation of the mathematics to justify certain mathematical moves, using such considerations to help determine where particular techniques and concepts may legitimately be used, thereby mitigating the risks of using unrigorous mathematics. Such reasoning does not fit neatly within existing versions of the mapping account, which focus on inferences from the mathematics to the world rather than inferences from the world to the mathematics.

Finally, I briefly show how to use an alternative account, the robustly inferential conception recently defended by McCullough-Benner (2019), to better represent the philosophically interesting features of these practices.

49. 'Fitting' a Genetically Engineered Mosquito into a Regulatory Schema and Scientific Risk Assessment Rubrics Meant for Chemicals: A Story of a Mismatch
Zahra Meghani

Genetically engineered (GE) organisms meant for release in the 'wild' present a challenge for regulatory agencies that function under the auspices of a biotechnology regulatory policy framework that did not anticipate their 'creation'. This presentation will discuss some of the problems with regulatory agencies attempting to regulate GE organisms using a categorization schema that fails to recognize them as living entities. To that end, the efforts by US regulatory agencies to regulate a GE mosquito intended to reduce the incidence of mosquito-borne diseases in humans will be used as a case study. The GE mosquito in question is a genetically modified Aedes aegypti mosquito, with a heritable synthetic genetic sequence that results in tetracycline dependency. That dependency causing trait, in effect, amounts to a lethality gene; the vast majority of the progeny of such insects cannot survive outside of laboratories unless they have access to tetracycline. Proponents of the use of GE insects intend to release sufficient numbers of male GE mosquitoes such that they will outcompete their male wild-type counterpart to mate with their female wild-type counterpart, resulting in offspring who will inherit the tetracycline dependency trait. It is expected that the larvae will not survive to adulthood because, presumably, adequate amounts of tetracycline will not be present in the regions where the GE mosquitoes are released. As a result, the population of Aedes aegypti in the target area will decline, presumably, resulting in lower incidence of transmission of Zika and dengue in humans. The GE insect was initially classified as a new animal drug, with the US Food and Drug Administration (FDA) asserting regulatory authority over it. The synthetic genetic construct introduced into the mosquito results in a change in its physiology; the FDA has authority over objects that cause such changes in organisms. In October 2017, the FDA ceded regulatory authority over the GE mosquito to the Environmental Protection Agency. The decision was made that the GE insect was a (bio)pesticide because it reduced the population of its wildtype counterpart, which is considered a pest.

Given that the GE mosquito is a living animal, the decision to construe it for regulatory purposes as a new animal drug or a (bio)pesticide raises questions about the kind and scope of the scientific risk assessments that have been and might be conducted by regulatory agencies to determine whether field trials of those GE organisms should be permitted. To address the issue of the quality of scientific risk assessments that construe a living animal as a new animal drug or as (bio)pesticide, this presentation will do the following:

i. Identify and evaluate the normative concerns shaping the decision to classify the GE Aedes aegypti mosquito as a new animal drug and then as a (bio)pesticide for regulatory purposes.
ii. Discuss the epistemic gaps and flaws that are likely to characterize scientific risk assessments that treat living organisms as chemicals.

50. Framing Cross-disciplinary Knowledge-generation Processes
Julie Mennes

Cross-disciplinary (CD) research appeared on the agenda of governments, research funders and universities in the 1970s (Klein, 1990). At the same time, cross-disciplinarity became itself a research subject in fields such as philosophy, sociology, education science, management science and scientometrics. The resulting literature focuses on the theoretical characterization of (different types of) cross-disciplinarity and the ways in which traditional institutional structures hinder collaboration across disciplinary boundaries (Mattila, 2005). An aspect that has not received enough attention are the processes by means of which knowledge is generated in CD research (O'Rourke et al., 2016). Yet, insight into these CD knowledge-generation processes is important for philosophers of science who want to keep track of contemporary scientific practice. It is also useful for researchers who are (interested in) performing CD research and want to reflect on the underlying knowledge-generation processes. Moreover, it is practically useful for research funders, who have to determine which research projects are cross-disciplinary and want to be able to compare CD research projects.

The (few) studies that are available tend to fall into two categories: they either (i) identify general, high-level features of CD knowledge-generation processes (e.g. "interdisciplinarity involves integration") or (ii) describe CD knowledge-generation processes in highly specific situations (e.g. "the combining of theories in sustainability science").

In my talk, I explore a middle way. I show how CD knowledge-generation processes can be described by a contextualist theory, i.e. a theory that can both (i) provide systematic insight into the similarities and differences between CD knowledge-generation processes in different CD research projects, and (ii) explicate the CD knowledge-generation processes that are at play in specific CD research projects. To do so, I use a case study of a CD project in urban studies.

The proposed contextualist theory consists of a structured set of explications of common concepts in the scientific and science studies literature. Two groups of concepts are important. The former are concepts that refer to actions and products distinguished in a discipline, such as 'data', 'model' or 'representation'; the latter consists of concepts that refer to interactions between disciplines such as 'bridging', 'integrating' or 'borrowing'. For the explication of these concepts, the theory makes use of conceptual structures that I will call 'frames'. A frame encompasses (i) couples, each of which comprises an attribute and its related 'potential value'-definition or list of value options and (ii) frame-internal logical relations between attributes, defining restrictions between attributes with respect to their value options. Roughly put, frames are used to capture explications of concepts that refer to actions and products, and concepts referring to interactions are explicated in terms of frame modifications.

The frames capture explications of general concepts but they can also be applied to a given CD research project by specifying them, i.e. by determining values of relevant attributes and selecting relevant frame-internal relations. Hence, frames are well suited to serve as the building blocks for a theory that allows to compare CD knowledge-generation processes across CD research projects and to gain insight into the CD knowledge-generation processes at play in a given project.

Klein, J. T. 1990. *Interdisciplinarity: History, theory, and practice*. Detroit: Wayne State University Press.
Mattila, E. 2005. Interdisciplinarity “In the Making”: Modeling Infectious Diseases. *Perspectives on Science*, 13(4).
O'Rourke, M., Crowley, S., & Gonnerman, C. 2016. On the Nature of Cross-disciplinary Integration: A Philosophical Framework. *Studies in History and Philosophy of Science Part C: Studies in History and Philosophy of Biological and Biomedical Sciences*, 56.

51. Scientific Progress in Scientific Practice: An Empirical Study
Moti Mizrahi

According to Chang (2007, 1), “Scientific progress remains one of the most significant issues in the philosophy of science today.” The aim of this paper is to contribute to the philosophical debate over the nature of scientific progress by taking an empirical approach. It reports the results of an empirical study of the following philosophical conceptions of scientific progress: the semantic account of scientific progress (progress in terms of truth), the epistemic account of scientific progress (progress in terms of knowledge), and the noetic account of scientific progress (progress in terms of understanding). For, as van Fraassen (1994, 184) puts it, “Any philosophical view of science is to be held accountable to actual scientific practice, scientific activity.” Accordingly, philosophical accounts of scientific progress should be held accountable to actual scientific practice, or scientific activity, as well. Philosophical accounts of scientific progress can be tested empirically against scientific practice by using the methods of data science and corpus linguistics. These methods allow for the study of large corpora of scientific texts systematically in order to uncover patterns of usage. In particular, do practicing scientists talk about scientific progress in their published work? If so, in what terms? Do they talk about scientific progress in terms of truth, knowledge, or understanding? Finding answers to these questions empirically might help shed new light on the nature of scientific progress by revealing how practicing scientists conceive of scientific progress. That way, philosophical accounts of scientific progress could be held accountable to scientific practice.

Using the methods of data science and corpus linguistics, databases of scientific publications can be mined for instances of the terms ‘truth’, ‘knowledge’, and ‘understanding’ in order to see whether, and to what extent, practicing scientists use these terms in their published works. If the semantic account is true, we should find practicing scientists using ‘truth’ more frequently than ‘knowledge’ or ‘understanding’ in scientific publications. If the epistemic account is true, we should find practicing scientists using ‘knowledge’ more frequently than ‘truth’ or ‘understanding’ in scientific publications. If the noetic account is true, we should find practicing scientists using ‘understanding’ more frequently than ‘knowledge’ or ‘truth’ in scientific publications.

In order to have a large and diverse sample that could be representative of science as a whole, systematic searches were conducted on data mined from the following subjects in the JSTOR database: Anthropology, Archeology, Astronomy, Biological Sciences, Economics, Geography, Geology, Linguistics, Mathematics, Paleontology, Physics, Psychology, Sociology, Statistics, and Zoology. That way, the datasets contain representative disciplines from the life sciences, physical sciences, and social sciences, as well as Mathematics and Statistics. The results suggest that, of the terms for the basic units of progress, namely, ‘truth’, ‘knowledge’, and ‘understanding’, it is ‘understanding’ that occurs more frequently than ‘truth’ or ‘knowledge’ in the context of talk about the aims or goals of science.

52. Internalizing Research Integrity: A Practice Based Project to Enhance Research Ethics Education

Barton Moffatt

Research integrity is central to successful epistemic practice in science. One consequence of the increase in research ethics regulation and the corresponding increase of research ethics bureaucracy is the unintended impression by some that research integrity is external to the practice of science. It is not unusual to hear researchers in certain disciplines refer to the bureaucracies of research ethics as the ethics police. This impression is unfortunate because people tend to defer ethical judgments if they perceive it to be someone else’s responsibility. Some research in economics suggests that external interventions may undermine intrinsic motivations in a crowding out effect (Frey & Jegen 2001). The more researchers perceive research ethics to be a domain external to scientific practice and adversarial in nature the less they will view ethics as necessary for their practice. Additionally, they will be less likely to embrace the mantle of teaching research ethics as an integral part of science.

This paper argues that there is a need to redirect this dynamic by educating graduate students in a way that highlights the centrality of research ethics to scientific practice. One way to do this is to introduce a moral exemplar project in research laboratories that makes students learn more about their broader research communities and identify scientists who have built a reputation for particularly moral behavior. Ideally, this will lead to a discussion within the laboratory that identifies why the behavior was important to the community. Of course, not all graduate students are well situated in active research communities; these students would have the option of taking a historical approach and researching scientific moral exemplars in the history of their fields. I argue that this type of project will enhance research ethics education by emphasizing the centrality of research ethics to epistemic communities and blunt the force of externalizing bureaucracies.

Frey, Bruno & Jegen, Reto. 2001. Motivation Crowding Theory: A Survey of Empirical Evidence. *J Econ Surv*. 15(5).

53. Intervention and Backfire in the Replication Crisis

Aydin Mohseni

Over the past decade it has been shown that a majority of findings in certain literatures in the social and biomedical sciences cannot be replicated, that many well-regarded findings in these literatures have been false, that journal prestige may be anti-correlated with replicability, and that failed replications of findings from preclinical trials cost industry and the public tens of billions of dollars per year in the United States alone. A host of problems have been imputed in producing the observed irreproducibility, and various interventions have been proposed to address these problems.

This paper addresses proposals for intervention in the replication crisis. I argue that prominent proposals neglect the social dimension of science, and that attending to the social dimension reveals insights necessary for effective intervention. I demonstrate---using the tools of game and decision theory and Bayesian inference---that accounting for heterogeneity in scientists' beliefs, incentives, and research practices reveals that remedial interventions can run afoul of unintended backfire effects. Changes in research practices and evidentiary standards that appear justified in the case of individual studies can be suboptimal at the level of science as a whole.

This work challenges standard accounts of scientific epistemology which formulate the norms statistical inference in terms of their consequences for the properties of individual studies, rather than on scientific practice as a whole. That is, accounts that cash out the norms of inference in terms of controlling error rates of statistical tests, of reliable detection of an error in the study hypothesis if one were present, or of yielding coherent predictive probabilities.

For example, consider two prominent proposals for intervention: lowering the conventional threshold for statistical significance from .05 to .005 and promoting the norm of preregistering studies. Each proposal is well-motivated from the perspective of its impact on the statistical properties of an individual study; yet one must also consider their effect on a population of studies.

I show that if studies using less sound methods (i.e., employing more questionable research practices such as p-hacking) are better able to attain statistical significance than those using more sound methods, then there will be a critical point beyond which further lowering the significance threshold will backfire, increasing the false discovery rate of the literature. In the case of voluntary preregistration, I show that if one of two conditions is met---(i) scientists whose study hypotheses are more likely to be true are also more likely to preregister their studies, or (ii) scientists who are more likely to be using more sound methods are more likely to preregister their studies---then preregistration can lead to an increase in the false discovery rate of the literature as a whole.

In each case, a failure to account for scientists' beliefs and propensities leads to misprescription of change in evidentiary standards and practices. Taken together, my arguments suggest that the norms for a better science must be sensitive to the characteristic challenges of the epistemic communities in question, and that any successful scientific epistemology must attend to the social dimension of science.

54. What Have We Learned About the Engram?
Jonathan Najenson

The engram, a hypothetical store in which information is held in the brain, is the linchpin of the sciences of memory. Until recently, the ability to find the engram was limited by the specificity of the tools available. The development of optogenetics, a new technology that enables the manipulation of neurons with light, is believed to have provided a principled way to locate the engram. This new technology provides to researchers unprecedented control that enables them to approach engrams in a way that was not conceivable before.

In my talk I look at how the use of optogenetic tools to answer questions about memory involves the application of different strategies to confirm discovery criteria. Neuroscientists offer three criteria to evaluate the discovery of an engram: contiguity, specificity and similarity. Each discovery criterion is confirmed by applying a unique strategy that makes use of optogenetic tools. I present these strategies and examine how they are used to determine the existence of engrams.

The first discovery criterion is contiguity. Following the acquisition of a memory, persistent changes resulting from the storage of learned information are expected to occur. To determine this, researchers employ a loss-of-function discovery strategy. Cells that were activated by a learning experience are labeled with a fluorescent protein and optical long-term depression (LTD) is later applied to these cells. Optical LTD reverses the memory consolidation process responsible for retaining encoded information, making these neurons unable to express their newly gained function. When amnesia of that experience results from such manipulation, researchers learn that without the relevant structure the memory could not be retrieved.

The second discovery criterion is specificity. A memory is believed to have a specific neural vehicle where it is retained. To determine if a neural population is specific to a memory, researchers employ a gain-of-function discovery strategy. In a gain-of-function experiment optical long-term potentiation (LTP) is induced in memory–specific cells that were previously labeled and functionally inhibited. In contrast to LTD, optically inducing LTP amounts to turning the memory on by artificially initiating memory consolidation. Using this strategy in an experimental setting with two different stimuli, applying optical LTP to these labelled cells results in behavior to only one of these stimuli, thereby confirming their specificity.

The last discovery criterion is similarity. The engram retaining the memory is expected to be structurally similar to the event that is remembered. The strategy employed here uses optogenetic tools to create or alter a memory by manipulating its neural vehicle. In this case, memories may be artificially created by associating them with learning episodes from another context. When stimuli are administered in a given setting, labeled neurons from another context are optically reactivated. The co-activation results in an artificial creation of a memory, connecting the original memory to the artificially activated context. By optically altering an existing memory in order to create an artificial one, researchers examine how changes to the structure of a neural vehicle alters the engram's content.

55. Changing Representations and Explanations of Proteins in the 20th century: The Conceptual Evolution of Allostery from Static to Dynamic
Jacob Neal

The discovery of allostery in 1961 led Jacques Monod to claim that he had 'discovered the second secret of life' (Ullman 1979, 167). This discovery revealed that protein functions could be controlled biochemically through feedback mechanisms. Allostery, or the phenomenon whereby binding at one site on a protein affects binding at a distant site, has become a central topic of research within structural biology over the past 60 years (Cui and Karplus 2008). Since it was first characterized by Monod, it has undergone a major conceptual evolution. Beginning in the 1960s, allostery was considered a structural phenomenon, and the explanations for it focused on the transition between discrete protein conformations that occurred upon binding a small molecule called an allosteric ligand (Monod et al. 1965). An alternate account developed in the 1980s redescribed allostery as a dynamic phenomenon. It explained allosteric behavior as resulting from changes in protein flexibility upon binding the allosteric ligand. This scientific episode is ripe for historical and philosophical analysis. At the heart of this story is a genuine historical puzzle. Although dynamic accounts of allostery were developed in the 1980s (Cooper and Dryden 1984), they gained little traction amongst the community of structural biologists in the 1980s and 1990s, most of whom were still working within the structural framework developed by Monod and his colleagues. It would take almost three decades before dynamic allostery would rise to prominence as an alternative account for explaining allostery. What explains the emergence of dynamic accounts of allostery and the long delay in their acceptance by the majority of protein scientists?

This paper offers a historical account that attempts to answer these questions. My analysis shows that dynamic accounts of allostery were developed in the 1980s by scientists with backgrounds in physics and biophysics who sought to discover the consequences of treating proteins as small thermodynamic systems. These novel accounts of this protein function were not developed in response to the empirical discovery of anomalous cases of allostery inexplicable by the Monod's structural model, nor were they developed in response to advances in technology. That is, they were neither anomaly-driven nor technology-driven. Instead, I argue that the development of dynamic accounts of allostery was primarily theory-driven: the goal of the scientists who developed these accounts was to offer explanations of allostery that considered the thermodynamic consequences of ligand binding on protein flexibility. The dynamic accounts gained a small following within the scientific community in the 1980s, and it was this handful of scientists committed to the thermodynamic treatment of proteins and protein functions that enabled dynamic accounts of allostery to arise as a rival to structural accounts. In the late 1980s and 1990s, these scientists—who were already committed to the dynamic framework—sought to discover empirical cases that could confirm their models, and it was the slow accumulation of these empirical cases of dynamic allostery that ultimately provided the necessary evidence for dynamic accounts of allostery to gain acceptance within the community of allosteric researchers in the 2000s.

56. Guaranteeing Objectivity in Animal Cognition Research
Eveli Neemre

In my presentation, I will distinguish and analyze the different objectivity strategies in animal cognition research. Animal cognition research is an interdisciplinary field where a variety of different methods are used to research the cognitive capabilities of various species. This research has been criticized due to some dubious research practices and concerns about anthropomorphism.

For example, Sebeok (2000, p. 50) criticized symbolic communication research by finding three main issues with these projects: „Inaccurate observations and/or recordings of ape behaviors; the over interpretation of ape performances; and the unintended modification of all animal`s behavior in the direction of the desired results “. All of these problems could be summarized as a loss of objectivity. Similar concerns have been raised about animal cognition research more generally. Whether concerning anthropomorphism (Andrews, 2015) or about methods (Andrews, 2016; de Waal, 2016).

Researchers in the animal cognition field are aware of the problems related to the field and have devised strategies to guarantee researches integrity and objectivity. Broadly, I distinguish two main strategies for objectivity in animal cognition research. Firstly, avoiding problematic assumptions and secondly, avoiding constructing a problematic environment for the research. The first strategy is related to theories and arguments. The main motivator here is to avoid assumptions that could be detrimental to research. The second strategy is related to methodology and conducting experiments. To get viable research results, it is important to construct experiments so that the animals researched get a chance to reveal themselves.

An example of avoiding problematic assumptions is de Waal´s (2016, p. 53) know-thy-animal rule: “Anyone who wishes to stress an alternative claim about an animal´s cognitive capacities either needs to familiarize him- or herself with the species in question or make a genuine effort to back his or her counterclaim with data.” Following this rule helps minimize unfounded speculations, false assumptions and implementing false knowledge. Bates (2018) echoes this rule by stressing that elephants must be studied as elephants, not as primates if we wish to learn something about their specific cognitive skills. The theoretical foundation of such rules is considered when devising experiments.

In my presentation, I will rely on the risk account of scientific objectivity to analyze examples of objectivity strategies in animal cognition research and show how they help guarantee objectivity in such research.

Andrews, K. 2016. Animal Cognition, *The Stanford Encyclopedia of Philosophy, Summer (May 6) 2016 Edition*.
Andrews, K. 2015. *The Animal Mind. An Introduction to the Philosophy of Animal Cognition*. Routledge.
Bates, Lucy A. 2018. Elephants – Studying Cognition in the African Savannah. In Nereida Bueno-Guerra and Federica Amici (eds.) *Field and Laboratory Methods in Animal cognition. A Comparative Guide*. Cambridge University Press, pp. 177-198.
de Waal, F. 2016. *Are We Smart Enough to Know How Smart Animals Are?* Granta.
Sebeok, T. A. 2000. *Essays in Semiotics – I Life Signs*. Legas.

57. What Is ‘Reading’ in Mindreading?
Nedah Nemati

In 1991, Jack Belliveau and colleagues reported the first use of magnetic resonance imaging to capture human brain function in real time: a structure-to-function mapping of the human visual cortex (Belliveau et al., 1991). Those colleagues reflect that Belliveau’s work was fueled by his vision for capturing human thought – a dream that almost thirty years later lives on, as investigators refine fMRI’s temporal and spatial resolution, add sophisticated statistical techniques, and shift to brain activity pattern-to-function mapping for greater accuracy. With enough data, powerful magnets, and the right theories for connecting patterns to function, we will read minds.

This ongoing refinement nonetheless assumes a certain framework concerning the basic activity of reading itself. Here, stepping back to Belliveau’s vision, I challenge a one-directional approach to mindreading. For this, I draw from Louise Rosenblatt’s reader-response theory as a framework for rethinking the ‘reading’ in mindreading as a transactional process (Rosenblatt, 1988).

Transactional reader-response theory holds that readers bring a wealth of background to a text and construct its meaning through this temporally and contextually sensitive engagement. Just as one might ask, “What do we bring to words we read off the page?” I ask, “What do we, as researchers, bring to concepts we ‘read’ off brain-generated fMRI signals?”

Toward answering the latter, I examine the practices of a recent study identifying neural correlates of abstract concepts, and identify three ways in which brain signals are made interpretable. The first involves examining participants’ assigned tasks, the second the shaping of concepts’ fluidity or fixedness, and the third an examination of statistical tools used to discover the concepts’ neural basis, focused on factor analysis. Interrogating what we bring to ‘reading’ images will better position us to distinguish meaning generated apart from the ‘read’ brain and to avoid taking for granted that it is all waiting to be uncovered, or read, by us.

Belliveau JW, Kennedy D, McKinstry R, et al. 1991. Functional mapping of the human visual cortex by magnetic resonance imaging. *Science*. 254:716–719.
Rosenblatt, Louise M. 1988. Center for the Study of Reading: A Reading Research and Education Center Report. *Writing and Reading: The Transactional Theory*. New York University: 416: 1-17.

58. Keeping Us Honest: Lie Detection, Polygraphs, Moral Neuroenhancement, and Integrity Scores

Jo Ann Oravec

The speech act of “lying” is often characterized as having significant religious, ethical, and practical importance, even though the basic notions of honesty and prevarication can be complex and hard to convey. This paper explores the research agendas and technological approaches that foster the development of lie detection and related cognitive engineering technologies from philosophy of science perspectives, with an emphasis on initiatives that use big data and artificial intelligence (AI) approaches. The process of lie detection has been construed as “use of a physiological measurement apparatus with the explicit aim of identifying when someone is lying. This typically comes with specific protocols for questioning the subject, and the output is graphically represented” (Bergers, 2018, p. 1). Some emerging forms of lie detection technologies incorporate the remote collection of data without notification of subjects as well as contain various AI- and machine learning-assisted analyses. Some initiatives take invasive approaches that include brain scanning and/or electrical brain modification, providing forms of proactive cognitive engineering. Accumulation of “integrity scores” or other ways of profiling individuals over time in terms of their propensity to lie is sometimes a part of these new research initiatives and technological development strategies. As discussed in this presentation, applications of the algorithms and methods involved may indeed have particularly negative outcomes for individuals whose cultural and demographic backgrounds inspire them to frame truth telling in ways that vary from the researchers’ and implementers’ assumptions; since these lie detection technologies are often used in wartime and international border crossing contexts, such cultural differences can be especially problematic. The paper analyzes how these lie detection and cognitive engineering efforts have been evaluated from a scientific perspective, often in ways that are readily challenged but that are deemed acceptable because of the perceived security and economic needs for the devices.

Although many polygraph and lie detection approaches are indeed restricted in use in certain legal settings, an assortment of new technologies that are labeled as “lie” or “cheating” detection have emerged that are being adopted in particular contexts (Oravec, 2018; Greenberg, 2019); this paper critically examines the appropriateness of this framing. Researchers are also introducing brain scanning as a way to detect lies and identify forms of mental concentration; the notion of “self-lie detection” has been investigated by researchers, with the projected potential for increasing personal insight about truthfulness through technological means (Echarte, 2019). This paper projects future technological developments and outlines the continuing need for ethical and professional vigilance on the part of researchers and system developers as they choose projects to work on and technologies to bring to market. The significance of truth telling in everyday life is expanding as societal attention to truth and falsity issues as well as social conformity increases. Research and development efforts on workplace and educational cheating and deception have also gained new dimensions in the advent of big data capabilities, and some of the resultant initiatives are in use today despite the fact that they are in the early stages of testing and evaluation. Deception-related behaviors have been construed as a continuing and somewhat vexing issue for many kinds of institutions as administrators increasingly impose metrics on evaluation processes and as greater shares of human interaction go online, so further investments in the development and use of lie detection and moral neuroenhancement systems are likely.

Balmer, A. 2018. *Lie Detection and the Law: Torture, Technology and Truth*. Routledge.

Bryant, P. (2018, December 21). Will Eye Scanning Technology Replace the Polygraph. *Government Technology*. Retrieved from http://www.govtech.com/public-safety/Will-Eye-Scanning-Technology-Replace-the-Polygraph.html

Darby, R. R., & Pascual-Leone, A. 2017. Moral Enhancement using Non-invasive Brain Stimulation. *Frontiers in Human Neuroscience*, 11, 77. DOI=10.3389/fnhum.2017.00077

Echarte, L. E. 2019. Self-lie Detection: New Challenges for Moral Neuroenhancement. In *Psychiatry and Neuroscience Update* (pp. 43-52). Springer, Cham.

Greenberg, A. 2019. Researchers Built an “Online Lie Detector.” Honestly, that Could be a Problem. *Wired*. Retrieved from https://www.wired.com/story/online-lie-detector-test-machine-learning/

Harris, M. 2018. An Eye-scanning Lie Detector is Forging a Dystopian Future. *Wired*. December 4. Retrieved from https://www.wired.com/story/eye-scanning-lie-detector-polygraph-forging-a-dystopian-future/

Maréchal, M. A., Cohn, A., Ugazio, G., & Ruff, C. C. 2017. Increasing Honesty in Humans with Noninvasive Brain Stimulation. *Proceedings of the National Academy of Sciences* 114(17).

Oravec, J. A. 2000. Interactive Toys and Children's Education: Strategies for Educators and Parents. *Childhood Education* 77(2).

Oravec, J. A. 2013. Gaming Google: Some Ethical Issues involving Online Reputation Management*. Journal of Business Ethics Education*, 10.

Oravec, J. A. 2018. Secrecy in Educational Practices: Enacting Nested Black Boxes in Cheating and Deception Detection Systems. *Secrecy and Society* 1(2).

59. Reconceptualizing Development in the Sex Differences Literature
Derek Oswick

There are several foci in the literature on sex differences in humans. Proponents of so-called 'hard-wired' sex differences often rely on various endocrine studies (ie girls with CAH behaving more 'tomboyish'; effects of prenatal testosterone exposure, etc), toy preference studies, test results for various cognitive skills (ie spatial reasoning, item rotation, etc) and the like. The emphasis in these studies has been on low-level, 'biological' explanations for a given phenomenon (ie hormonal 'shaping' of brains, etc). Critics, both neurofeminist and otherwise, have responded largely with methodological critiques, as well as analysis of the types of sexist background assumptions that contribute to overlooking methodological flaws. Alternative explanations, where offered, have understandably highlighted social causes of differences. For example, stereotype threat analysis draws on the ways in which a test subject conceives of themselves as belonging to certain social groups (and the expectations/connotations that such membership brings), which then can increase cognitive stress and lead to test results that are in keeping with what one might expect from a member of the group in question.

My concern in this paper is that the way in which the respective sides have engaged with each other seems to resemble a nature vs nurture divide. (To be clear, I do not take the neurofeminists to be advocating a staunch nurturist position, but the emphasis on social factors does lend that impression). Certainly the way in which the material is discussed in the public forum lends credence to this impression; charges on Twitter of 'blank slatism' abound against the critics, who are often quick to fire back about 'biological determinism'. There is a clear ideological difference between the groups, and this type of engagement is prone to go back and forth as new 'landmark' studies highlight some new genetic driver of a given behaviour, vice versa for socialization. This sort of constant back and forth does not seem ideal for resolving the issues.

In this paper I will argue that to move beyond the current state of the literature we need to reconceptualize how we understand development in humans. Specifically, to dissolve the nature/nurture binary and emerge from under its long-reaching shadow, we need to understand development as an emergent property of constant interactions of various kinds within the ontogenetic niche. As Fine has put it, we need to understand development as 'development from' some prior set of interactions, rather than development 'towards' some fixed and predestined end point. The ontogenetic niche literature helps us keep in mind that there are multiple systems of inheritance, and so this type of approach lets us integrate both genetic and exogenetic studies since a developmental niche will never have genes without an environment, and vice versa.

60. Building ROBOT for Data Journeys
James A. Overton

Sabina Leonelli's "data journeys" call attention to the many steps and stages, path dependencies, diverse participants, and other crucial details that connect field measurements to scientific publications, databases, reports, and points between. Many data journeys in modern science are so long and complex that no single person can hope to understand every detail, raising difficult questions about epistemology, responsibility, and accountability. Contributing to that complexity are software and information systems that make much modern science possible, but require technical expertise that few scientists can claim. Many data journeys also include scientific ontologies, which are both terminology systems and technical artifacts built to standardize communication among people and machines.

In this paper I apply Leonelli's data journey approach to my first-hand experience as the lead developer of ROBOT. ROBOT is open source software that helps automate important tasks for ontology developers, such as managing large sets of terms, importing terms from other ontologies, automated reasoning, quality control, and release. It is used in dozens of Open Biological and Biomedical Ontology (OBO) projects, such as the Gene Ontology and the Ontology for Biomedical Investigations. As such, ROBOT plays at least a small part in a great many data journeys.

First I discuss how ROBOT has affected the OBO community since it was introduced five years ago. ROBOT has made technically-skilled ontology developers much more productive, allowing them to work on a larger number of larger ontology projects. It allows scientists with some technical skills to perform certain new and complex tasks themselves. Unfortunately there are also cases where adopting ROBOT has made it more difficult for ontology developers not familiar with ROBOT to continue contributing to those projects in the direct ways that they contributed before. At a larger scale, ROBOT is designed to make OBO best-practises easy, but also “bakes them in” to software that (despite being open source) only a small group has the technical skills to modify.

Second I reflect on three design principles that guide ROBOT development and their connection to Leonelli's concerns about distributed expertise and responsibility in data journeys. One is that, when applicable, ROBOT automation should behave like Protege, a desktop application that predates ROBOT and that most ontology developers use for manually editing ontologies. Another is that ROBOT should only add new functionality, never altering existing functionality. This is in tension with the need to fix bugs, and between the software's behaviour as intended, as documented, and as implemented. A third principle is that all new features be driven by multiple, real-world use cases. These design principles help build and maintain a smooth connection between ROBOT, user expectations, and the many different situations of use.

I argue that the data journey approach elucidates these design principles, while my case study extends Leonelli's work on scientific ontologies to recent advances in automation. Both demonstrate the “participative, reflexive management of data practices” that she endorses.

61. A Valuable Approach to Complexity: Interest-Driven Model Choice in Paleoclimatology

Meghan Page and Monica Morrison

This paper develops a novel way that values play a role in climate science through a discussion of paleoclimatic reconstructions of the mass-extinction event at the Permian Triassic Boundary (PTB). The PTB extinction is of growing interest to current scientists because it was presumably initiated by a significant increase in CO2 levels which triggered a warming event. Gaining insight into the trigger and trajectory of the PTB extinction may offer relevant insight into how current warming trends will impact various earth systems.

However, due to computing constraints, there is no “best” approach to modeling conditions at the PTB. Scientist must choose between complex models that represent numerous small-scale processes for a short period (AOGCMs), or low-resolution, simplified models that reconstruct a few particular large-scale processes for longer periods (EMICs).

AOGCMs are complex models that couple together various components of the earth system. They operate at a high resolution, allowing accurate simulations of current climate trajectories by resolving small-scale climate processes. However, because AOGCMs are so complex, they are very costly to run and require a large amount of processing time, so they are only used to reconstruct relatively short time periods (about 3000 years maximum). This can be problematic when reconstructing an extinction event that may have spanned a million years. To run simulations over longer time periods, scientists turn to less complex models known as EMICs. EMICs are coupled models that include fewer processes and run at a lower resolution than AOGCMs, meaning they ignore many small-scale processes. Because of their simplicity, EMICs are less successful at reconstructing the global climate state but are quite effective at modeling the evolution of particular processes over long periods of time. However, EMICS model individual processes in relative isolation, so they don’t take into account how these processes are affected by many of the other changes to the climate state.

Scientists must choose whether to generate a more holistic simulation of Permian climate that only reflects a short period in time or examine an extended evolution that focuses on one (or several) particular processes but ignores many contributors to the climate state. In the case of model choice, there is no clear scientific or epistemic reason to prefer one model or the other. Hence, which model a scientist chooses to use is largely determined by the kind of knowledge valued by the scientist, and this choice is based on values. However, we argue this appeal to values is legitimate, because it drives scientific pluralism, which is required for the study of complex phenomena. Following Helen Longino, we argue that complex phenomena, such as the climate, must be investigated from multiple competing and sometimes contradictory angles in order for a holistic picture to emerge. The differing values of scientists, then, ultimately enhance scientific objectivity as they propel scientists to investigate a single phenomenon from a diversity of perspectives.

62. Nuclear Modeling and Process Realism: What’s Really Going On
William Penn

Nuclear models are incompatible in their thing-terms: terms that refer to static entities like objects, structures, and substances. Specifically, the two most prevalent models: the liquid drop and shell models, treat the nucleus, its internal structure, and the component nucleons as entities that contradict each other’s properties. These differences allow these two models to describe and explain different nuclear experiments: fission and scattering in the liquid drop model, and single-nucleon excitation and nuclear decay in the shell model. However, prima facie, these differences also suggest that these models are incompatible in their ontology. Indeed, by maintaining adherence to standard static ontologies, henceforth “thing realism,” this incompatibility is irresolvable. That said, by taking seriously the experimental methods by which these models are constructed and the calculational tools these models provide for interpreting experimental outcomes, I construct a new form of realism about these models that renders them ontologically compatible. Namely, I argue that nuclear models are consistent in the dynamic entities to which they refer. Therefore, I advocate a pure process realism with respect to nuclear models. Critical to this process realism is the recognition that the processes referred to within nuclear models are essential parts of the observational acts that form nuclear experiments. In particular, because the dynamics within the nucleus must always be a continuous intermediary of our experimental interventions and the receptions of signals from the system, these dynamics are essential dynamic parts of nuclear experiments. We are therefore licensed in inferring these dynamic parts on the basis of experimental practice alone. In contrast, the thing terms reified by the thing realist in these models require additional inferences, the premises of which cannot be supported on the basis of experiment alone. The essential premise of these additional inferences is one of two options (a) that all dynamics (metaphysically) require statics to underlie them, or (b) that the existence of stability in an experimental system necessitates something static and unchanging within the system. These premises are question-begging if deductively supported, and insufficient if inductively supported. Thus, inferences to processes are experimentally supportable, whereas inferences to things are dubious at best. Process realism is therefore superior to thing realism in the context of nuclear models because it (1) establishes cross-model consistency, (2) accords with experimental practice, and (3) is epistemically modest.

63. Molecules, Scientific Realism, Orbitals, and Understanding
Myron A Penner and Amanda J Nichols

Chemists have represented molecules as three-dimensional objects with a particular orientation in space for over a century. Throughout that period and into the present, many chemists have thought about molecular structures in realist terms. That is, many chemists have thought and continue to think that the three-dimensional models of molecular structures are reasonably accurate depictions of how atoms and subatomic particles bond, form molecules, and how they actually are oriented in physical space. This is so, even though (for the most part) we don't have photographic images of molecular microstructures and subatomic particles. In our paper, we advance a novel argument for molecular realism based on molecular symmetry and IR spectroscopy. We then consider objections to molecular realism based on the physics of quantum field theory and the philosophy of scientific understanding.

Chemist's understanding of molecular symmetry is a very powerful example of how armchair theory and laboratory observation can combine to advance scientific understanding. With a small set of assumptions about how particular elements bond, one can create a model of molecular structure. As a result of a model's particular spatial orientation, one can perform symmetry operations (e.g. rotating the model 180 degrees, etc.) on the model and categorize molecules into point groups based on their symmetry elements. But while molecular structures might be modeled and grouped from the armchair, these groupings make testable predictions. For example, we can predict how a molecule will respond to IR spectroscopy based on its symmetry elements. Moreover, we can predict that molecules in the same point group (even though their physical structures look quite different) will respond to IR spectroscopy in similar ways. Using group theory mathematics, we demonstrate that both these predictions are confirmed, thus providing strong evidence that the armchair models of molecular structure are accurate.

Quantum field theory can provide strong objection to thinking that our molecular models are true or approximately true representations of molecular structures. In order to incorporate quantum mechanics, concepts of molecular orbitals have become part of the modern molecular structure theory, and this has led some to object to entity realism at the subatomic level. We argue that entity realism doesn't require realism about orbitals, just about the electrons that inhabit the orbital space. We also consider objections to molecular realism based on the claim that while molecular models do contribute to understanding, understanding doesn't entail realism. That is, although we draw a tight connection between representations of molecular structure, understanding of their chemical properties, and realism, achieving scientific understanding doesn't necessarily entail scientific realism (de Regt, 2017). Perhaps the way chemists represent molecular structure is best seen as an idealization that functions as a "felicitous falsehood" (Elgin, 2017). However, we argue that molecular realism has both theoretical and pragmatic advantages over considering these models as false idealisations. As such, while achieving understanding doesn't necessarily entail realism, we argue that there are good reasons to show that in this case, it does.

64. Physical Magnitudes and Physical Concepts: How Viscosity Challenges the Mapping Account of Mathematical Representation
Stephen Perry

One of the more popular approaches to articulating the role of mathematics in scientific modeling and explanation has been what is called the "mapping account." The mapping account supposes that there is some isomorphism or homomorphism between a mathematical representation and the physical phenomenon it is representing. A notable recent formulation of the mapping account is given by Christopher Pincock (Pincock, 2007b, 2012). In this formulation, Pincock introduces the notion of a matching model in order to accommodate infinite idealizations (Pincock, 2007a). While this account is an improvement on the "naive" mapping model, I will argue that Pincock's account and all mapping accounts rely on a correspondence between a mathematical structure and some structure in the physical world. Pincock's matching models address certain concerns about this correspondence critique, but I will argue that they still rely on a correspondence between the physical magnitudes represented in mathematical equations, and the world. I will use the case of representing viscosity in the Navier-Stokes Theorems and Prandtl's Boundary Layer Theory to challenge this notion of correspondence between a physical magnitude and the world.

To challenge the mapping accounts' correspondence requirements on physical magnitudes, I will investigate the concept of viscosity, particularly on the mapping account interpretation of these equations. Building on Morrison's (Morgan & Morrison, 1999) analysis of Boundary Layer Theory, I will argue that the implementation of viscosity in Navier-Stokes and Boundary Layer Theory involves a balancing between macro- and microlevel phenomena. The derivation of the Navier-Stokes equation involves both molecular assumptions about viscosity and macro-level intuitions, particularly in Navier's choice to move from a molecular to continuum-mechanical treatment. This move is an example of infinite idealization. I argue that the mapping account interpretation of Navier-Stokes fails to adequately represent these mechanics of how the viscosity concept is functioning in the Navier-Stokes equation, which leads to the correspondence problems the mapping account runs into when trying to reconcile Navier- Stokes and Boundary Layer Theory. In particular, I will use Wilson's (Wilson, 2006) and Chang's (Chang, 2007)accounts of how scientific concepts evolve in physical theory. Using a case study on the concept of force, Wilson discusses how the historical development of a concept affects its current use. Similarly, Chang discusses the problems in coordinating theoretical conceptions of temperature and measurements of temperature. While these accounts differ, they both recognize the fine grain resulting form the historical development of these concepts. These details do not enter into the mapping account, and this is what results in the correspondence problem that I identify.

65. A Framework for Analyzing Engaged Philosophy of Science
Kathryn Plaisance and Kevin Elliott

The last decade has seen a significant increase in 'socially relevant' and 'engaged' philosophy of science. This is reflected in the introduction of concepts such as 'field philosophy' and the development of organizations like the Joint Caucus of Socially Engaged Philosophers and Historians of Science (JCSEPHS) and the International Consortium of Socially Relevant Philosophy of/in Science and Engineering (SRPoiSE). Despite the increasing prevalence of engaged work — and growing attention to particular examples of it in the philosophical literature — there has been little in-depth analysis of the wide variety of approaches philosophers use, the particular audiences with which they engage, and the barriers to engagement. In this talk, we systematically examine the nature of such engagement.

We begin by providing some historical perspective. We note that previous accounts of engaged scholarship excel at motivating the value of having philosophers of science reach out beyond their discipline, but they tend to focus on particular ways of being engaged rather than providing a comprehensive account of different approaches to engagement. Building on these accounts, as well as on literature from other disciplines on the topic of engagement, we provide a general framework for characterizing different forms of engaged philosophy of science. Our framework focuses on two key dimensions of engagement: the social/communicative and the cognitive/epistemic. Both of these dimensions lie on a spectrum, meaning that engagement itself can come in a matter of degrees. After presenting the framework, we illustrate it using several exemplars of engaged philosophy of science. Finally, we discuss two ways this framework can guide future philosophical work. First, it can help those interested in engagement to pursue approaches that are best suited to their goals and contexts. Second, it can enhance future scholarship on engaged work by showing how the aims, partners, barriers, and outcomes of engagement can vary depending on where they are situated in this framework.

66. Opening up HPS-debates: On Reading Kuhn and the History of the Quantum
Jan Potters

Since a few decades, historical studies of the early quantum (Darrigol 2001, Gearhart 2002) have taken a meta-historical turn: they are not only interested in how Planck saw the discrete energy elements he introduced, but also in how historians in the past have studied this question. As is pointed out in these studies, this turn is a consequence of Kuhn's work on black-body theory, in which it is argued that Planck, up until 1906, understood his work not in terms of a quantized discontinuity, but rather in classical, continuous terms. These authors argue, more in particular, that different positions with respect to Planck's work are possible, depending on the precise interests that one has.

In this talk, I will argue that this meta-historical debate is based on a theory-driven reading of Kuhn's work, and that, following Rouse (1987)'s practice-focused characterization of

Kuhn's notion of paradigms and revolutions, a different reading of Kuhn's work is also possible. On this reading, Planck's work on the quantum pre-1906 should be characterized primarily as an approach that allowed him to derive solutions to certain problems and new contributions to the foundations of physics, without presuming that Planck already had a fully worked out understanding of the meaning of his energy elements. By means of a discussion of how Planck's work was received and discussed between 1901 (when he presented his energy elements) and 1911 (when it was the subject of the first Solvay conference), I will then argue that this meaning was rather elaborated through different theoretical and experimental investigations of how Planck's approach could be applied to other questions.

On the basis of this, I will then argue that the meta-historical debate can be enriched by means of this practice-focused reading of Kuhn and the history of the early quantum. Moreover, it can also enrich philosophical discussion on practices, as I will then show by means of a discussion of how perspectivists such as Giere (2013) have used Kuhn's work to elaborate their notion of a scientific perspective. This will lead me to argue that, as it stands, the perspectivist's conceptualization of scientific practice is also too theory-focused, and on the basis of my discussion of Kuhn and the quantum history, I will then sketch, as a sort of conclusion, how a more practice-focused understanding of perspectives could look like.

Darrigol, O. 2001. The Historian's Disagreement over the Meaning of Planck's Quantum. *Centaurus* 43.
Gearhart, C. 2002. Planck, the Quantum, and the Historians. *Physics in Perspective* 4.
Giere, R. 2013. Kuhn as Perspectival Realist. *Topoi* 32.
Rouse, J. 1987. *Knowledge and Power*. Cornell University Press.

67. Cloudy with a Chance of Constraint Tuning
Justin Price

There are two discussions in philosophy of science about modeling that this paper seeks to combine into a new perspective on the epistemology of science. One discussion is about the transfer of models across scientific disciplines. Certain models seem capable of transfer across scientific domains while others do not. Philosophical accounts for model transfer identify the features that transferable models share – a mathematical core for instance – and explain why these features drive transfer, affording less computational demanding methods for mathematical modeling. Research so far provides a framework for investigating the scientific practice of model transfer, revealing unique modes of conceptual strategizing in the sciences.

The other discussion is about the role of idealizing in scientific modeling. Idealizations are distorting features of models that seemingly involve misrepresentations or falsifications of phenomena. This conflicts with the face-value role of scientific models as accurately and truthfully depicting nature. If idealizations are in tension with a model's accuracy or truthfulness, then it is puzzling that scientists so often choose to idealize. Philosophers of science resolve this tension by theorizing epistemic aims like non-factive explanation or understanding as what idealizing contributes towards.

This paper offers a smaller scale alternative, made salient by analysis of a case of model transfer, for the epistemic role of idealizing: constraint tuning. Constraint tuning is new concept I introduce that identifies the practice of delimiting the area of application of a mathematical model through idealizing. It is a method that scientists use to forward their modal knowledge of a target system by making explicit what counts as a representation. This process involves thoughtful conceptualization and drawing up formal mathematical constraints, and fits in a developing philosophical discussion on conceptual strategizing in the sciences.

To illustrate how constraint tuning is another epistemic aim of idealizing, I introduce a new case study on the transfer of the virial theorem to astrophysics in the 1960s. Model transfer, instead of being for reasons to do with tractability, involves affording new modal knowledge to a domain. The transferred model does so by allowing scientists to delimit a new region for modal exploration. In this case, the transferrable model core - the virial theorem - consists of an equation that modelers use as a formal mathematical constraint for drawing up principled idealizations. In astrophysics, the introduction of the virial theorem allowed scientists to explore the supposition that interstellar clouds were essentially stable things – an idealization. This transfer of the virial theorem then allowed other modelers to piggyback on this supposition, using this idealization in models of stellar evolution, generating hypotheses and inferences about the time scale for the birth of stars. What modelers introduced to astrophysics when they began using the virial theorem was a new way to delimit which things in model systems were stable clouds. This allows scientists to explore, by varying initial conditions and other parameters, how clouds may possibly evolve.

68. Investigating the Grammar of Sustainability
Jeff Ramsey

It is widely recognized and accepted that humans are having a negative and perhaps irreversible impact on the global environment. The impacts affect not just the environment ‘out there’ but also our livelihoods, our health, our economies and our quality of life. Sustainability has been offered as a way to stop the decline and create new ways of good living that can continue for longer periods of time and, with any luck, indefinitely.

How can we use the notion of sustainability to get us to a better place? One widespread, almost ubiquitous strategy is to offer a theory of sustainability. Current theories are couched in a wide variety of bases: biophysical theory; coupled socio-ecological systems (SES); various forms of economics; trans- and interdisciplinary science; science-and-technology-studies (STS); ethics and values; and various combinations of these bases. Generally speaking, the hope is to generate a basis which will address what sustainability means normatively, conceptually and practically.

I argue that many such attempts to theorize sustainability aim for, echoing Wittgenstein in the *Philosophical Investigations*, “complete exactness” (PI §91). In attempting to produce a theory of sustainability, they adopt an analytical approach that aims to assign strong, referentially determinate core meanings to the various elements of sustainability. This is revealed by their sometimes explicit but often implicit use of reductionistic, linear, naïve empiricist and compositional strategies to theory building. I substantiate this claim through analysis of prominent sustainability proposals such as Rockström’s planetary boundary approach; Matson, Clark and Andersson’s *Pursuing Sustainability* (Princeton, 2016), the work of Bryan Norton, and Curren and Metzger’s *Living Well Now and In the Future* (2017).

I argue that such strategies push us into unhelpful and unjustifiable pictures of what sustainability is and can be. They attempt to characterize sustainability independently of its social context as practiced by particular groups in particular places. As an alternative, I argue that we should aim for “complete clarity” (PI §133). A better approach to a ‘theory’ of sustainability involves seeing theories as: 1. via John Norton and Alan Love, principles linked via material (rather than formal) inferences, and 2. via Mark Wilson, a patchwork of projected and extended meanings for scientific terms (as opposed to a formal structure of already given, statable-in-advance stipulations). This approach leads to a descriptively (and normatively) thick characterization of sustainability. Since most environmental problems are ‘wicked,’ with causes and effects at multiple scales and with the problems changing as we intervene in order to address them, such a view better provides better ways of thinking about the theoretical commitments of sustainability proposals. It offers a way of thinking about the grammar – rather than the logic – of sustainability.

69. Does Species Status Support Conservation Status?
Thomas Reydon

The IUCN Red List of Threatened Species is perhaps the most widely known tool in nature conservation. It currently contains more than 112,400 species of animals, plants and fungi, ranked according to seven threat levels. The Red List presents itself as “the world’s most comprehensive information source on the global conservation status of animal, fungus and plant species” intended as a ready-for-use tool for conservation practice that “shows us where and what actions need to be taken to save the building blocks of nature from extinction” (see https://www.iucnredlist.org/about/background-history).

An implicit assumption underlying the Red List is that there is something special about the species rank that distinguishes it from the other ranks in the Linnaean hierarchy and merits placing emphasis on species in conservation efforts rather than on families, genera, subspecies, and so on. This assumption is expressed by highlighting species as “the building blocks of nature” (see above). It is implicitly endorsed by many conservation organizations that have species conservation as one of their focal aims, such as the World Wildlife Fund, the National Wildlife Federation, and others. But does the fact that a group of organisms is classified as a species support its being emphasized as a unit of concern in nature conservation?

My answer to this question will be partly affirmative, opposing a skepticism about species conservation that stems from the persistent problem of determining the nature of species. The extensive literature in the philosophy of biology on the species problem clearly shows that thinking of species as in any straightforward sense being building blocks of nature is mistaken. Also, the literature shows that there is no unique way of carving up the living world into species: some species concepts yield more than twice as many species than other concepts when used as the basis for classifying the organisms in a particular area. This gives serious reasons to assume antirealism about species and to doubt that species should stand at the focus of conservation efforts. While agreeing with this line of reasoning, I will maintain that ranking a group as a species does support its status as an important focal unit in conservation in ways that ranking a group as a subspecies, genus, etc. does not. My positive proposal will be that species should not be thought of as constituting a particular category of kinds or entities, but rather being a species should be thought of as a status attributed to a group of organisms on the basis of its role in the evolutionary process.

I will argue that this view of species as changing parts of an ongoing process (rather than static building blocks of the living world) allows the attribution of species status to a group of organisms to support normative claims in conservation practice in a way that is compatible with antirealism about species.

This talk will show how in the case of species theoretical research in philosophy of science bears strong relevance for an area of practice (i.e., nature conservation).

70. Representing as Reasoning: Reimagining Inferentialism
Mark Risjord, Kareem Khalifa and Jared Millson

In the scientific representation literature, most parties agree that scientific representations must support surrogative inference, in which scientists reason from models to their respective targets. “Substantivists” seek to explain models’ inferential roles in terms of mapping relations such as denotation, similarity, or some kind of morphism. By contrast, “inferentialists” seek to explain most other aspects of modeling in terms of models’ capacity to support surrogative inferences. Inferentialists are frequently charged with rendering the concept of surrogative inference mysterious, and thereby covertly presupposing substantive mapping relations. Call this the opacity problem. In this paper, we solve the opacity problem by offering a richer account of surrogative inference, and then showing that it renders appeal to substantive mapping relations unnecessary. In this way, we offer a new inferentialist conception of scientific representation that provides sufficient conditions on scientific representation without appealing to a substantive relations.

We begin by describing the opacity problem in greater detail. Roughly, inferentialists deny the philosophical importance of finding substantive relations that explain how models enable surrogative reasoning. However, most accounts of inference presuppose substantive relations. Consequently, inferentialists are either covertly appealing to these substantive relations, or else are making surrogative reasoning more mysterious than it actually is. So construed, the opacity problem poses a significant challenge for the two leading inferentialist accounts—Suárez’s and Hughes’.

Solving the opacity problem requires inferentialists to describe surrogative inference in greater detail than has been done thus far. To that end, we use examples of social-psychological and kinematic models to argue that surrogative inference is justified when it arises in the right way from the use of the model in a particular context of inquiry, what we will call the inference’s inferential pedigree. The inferential pedigree requires the investigator to secure three claims: (1) that the possible derivations from the model would be answers to the question(s) guiding the inquiry, (2) that elements of the model have been specified by appropriate operationalization and reliable measurement, and (3) that the specified model entails a premise of the surrogative inference.

With this more detailed account of how surrogative inference is justified in hand, we then argue that if a user is entitled (by having an adequate inferential pedigree) to use model M to reason surrogatively about a target T, then M represents T (for that user). Specifically, we first argue that inferential pedigrees do not presuppose any substantive relationships, such as denotation, similarity, or morphism. We then rebut various challenges purporting to show that some further substantive relationship must be added to a model’s inferential pedigree for it to represent its respective target. By showing that none of the major candidates for substantive relationships are required, we conclude that surrogative inference is sufficient for scientific representation.

71. Sex is a Biosocial Kind

Esther Rosario

Since the 1970s Anglo American feminists have espoused the slogan "gender is the social meaning of sex." Sally Haslanger notes that this slogan has taken on a variety of interpretations, but each interpretation tends to lend itself to the view that gender is social. A natural question to ask in turn is whether sex is biological. In this paper, I shall answer this question in the affirmative, but with an important qualification. I'll argue that sex is biological, but given the ways that sex is categorized and practiced, it also has social features. I shall maintain that sex is a biosocial rather than a natural kind. In order to motivate this claim, I'll consider Judith Butler's strong social constructionist account of sex and Muhammad Ali Khalidi's natural kinds approach, which takes an evolutionary perspective. I shall contend that Butler's view, although ground breaking, falls short of offering a scientifically informed account of the nature of sex. Her view overlooks the overwhelming evidence that sex has a biological basis. Further, I'll hold that Khalidi's view, while scientifically informed, is too narrow in scope. His view only focuses on gametic sex and behaviour when there are a variety of sex-based features and relations that are explanatorily significant. I'll maintain that what is needed is an account that first analyses sex from a developmental perspective, then tracks how sex is imbued with social structures and practices. It is crucial to assess not only how sex is classified (epistemically), but also whether or how social structures and practices (metaphysically) construct social individuals with sex-based properties. I'll maintain that sex-based properties are properties that include genes, hormones, gonads, genitals, and other secondary physical features. Sex-based properties vary along continuous axes in ways that are generalizable but do not provide necessary and sufficient conditions for 'female' or 'male' group membership. I shall hold that the idea that individuals have sex-based properties is more advantageous than the notion that individuals are either 'female' or 'male.' The reason for this is that individuals can have various combinations of sex-based properties, even for those who have a typical 'female' or 'male' appearance. Although social individuals include other animals, my account will focus on humans. In order to highlight the ways in which humans have sex-based properties that are imbued with social meanings, I shall draw on the case of the medical management of intersex bodies. The existence of intersex individuals is evidence that sex is not a binary biological category. However, intersex bodies are sometimes physically altered through unnecessary surgeries and hormone treatments without the consent of the individual. This occurs for the purposes of normalizing the appearance of intersex bodies. I'll contend that social meanings about what it is to be 'female' or 'male' help constitute how sex functions as a kind. I'll maintain that it is critical to acknowledge that sex is both biological and social because how we categorize sex in the social world influences the sex-based properties and appearance of individuals. Only viewing sex as biological ignores the very real ways that sex is socially practiced, and forecloses the possibility to change those practices.

72. The Extended Evolutionary Synthesis and Laws-in-Scientific-Practice

Joseph Rouse

Recent controversies in evolutionary biology provide an extraordinary opportunity for understanding conceptual or theoretical conflict and change in the sciences. Multiple challenges to the Modern Evolutionary Synthesis (MES) developed over four decades are now controversially consolidated by advocates of an Extended Evolutionary Synthesis (EES) (e.g., Pigliucci and Müller 2010, Laland et al. 2015, Müller 2017), and criticized by MES advocates (e.g., Futuyma 2017). Revisionist conceptions of scientific lawfulness developed by Lange (2000, 2007), Haugeland (1998) and Rouse (2015) provide new insights into the shape and significance of these disputes.

Both the MES and the EES developed over multiple decades, with important earlier antecedents. Both span multiple biological sub-fields, which often implement their over-arching theoretical standpoints with distinctive emphases. EES advocates accept core principles of the MES, while limiting their explanatory scope and range of inductive projectibility. Advocates of both syntheses mostly accept the same empirical results, while differing on their theoretical significance, the explanatory resources they require, and the explanatory aims. MES defenders argue that the EES mixes phenomena already known and traditionally understood with unsupported speculation. EES advocates argue that the explanatory principles guiding standard explications of evolutionary theory block or marginalize inquiry into important aspects of evolution, mostly via an unfounded explanatory priority for only some contributing causes of complex phenomena.

Lange, Haugeland, and Rouse reject philosophically pre-conceived notions of scientific laws. Conceptions of laws and lawfulness should instead be governed by laws' roles in scientific practice, for inductive confirmation, explanation, and counter-factual reasoning. Laws are norms of reasoning within a conceptual domain rather than always-invariant truths. Laws differ from accidents not singly, but as open-ended sets whose range of counterfactual invariance is defined with respect to the set. Sets of laws function in these roles only when conjoined with normative governance of their relevant counterfactual scope, the skills and practices with which their concepts are applied, and the conditions under which the laws hold, including ceteris paribus qualifications and noise-tolerance.
Recent presentations of the EES/MES disagreements focus on four primary loci: sources of variation, significance of developmental plasticity, scope of inheritance, and evolutionary significance of niche construction. These disagreements primarily concern lawfulness in this sense: the scope of the evolutionary domain (especially concerning the place of developmental processes and ecology in evolution) and the normative considerations that govern inferential relations. Concluding considerations address the suggestion, made constructively by Godfrey-Smith (2001) about an EES precursor and critically by some MES defenders, that differences between syntheses are more philosophical than scientific.

Futuyma, D. 2017. Evolutionary Biology Today and the Call for an Extended Synthesis. *Interface Focus* 7.

Godfrey-Smith, P. 2001. On the Status and Explanatory Structure of Developmental Systems Theory. In Oyama, Griffiths, and Gray, eds., *Cycles of Contingency*, ch. 20. MIT Press.

Haugeland, J. 1998. *Having Thought*, ch. 13. Harvard University Press.

Laland, K., Uller, T., et al. The Extended Evolutionary Synthesis: Its Structure, Assumptions, and Predictions. *Proc. R. Soc. B* 282.
Lange, M. 2000. *Laws in Scientific Practice*. Oxford University Press.
Müller, G. 2017. Why an Extended Evolutionary Synthesis is Necessary. *Interface Focus* 7.
Pigliucci, M. and Müller, G., ed., *Evolution: The Extended Synthesis*. MIT Press.
Rouse, Joseph 2015. *Articulating the World*, ch. 8-10. University of Chicago Press.

73. Is a Bird in the Hand Worth even one in the Bush
Carlos Santana

As heated debate over the past couple of decades goes to show, the titular question has far from a straightforward answer. At issue is the traditional practice of collecting voucher and type specimens for nomenclatural and taxonomic purposes, particularly in the case of vulnerable populations and endangered species of vertebrates. Many researchers have argued that photography and non-lethally collected tissue and genetic samples can be sufficient for scientific purposes (e.g. Donegan 2008; Minteer et al. 2014; Pape 2016). Others, most prominently researchers working in museums, have taken to the pages of scientific journals defending the traditional practice of collecting complete specimens (e.g. Cotterill 1997; Bates et al. 2004; Dubois and Nemesio 2007; Rocha et al. 2014; Ceríaco et al. 2016).

Havstad (2019) convincingly demonstrates that there are good considerations on both sides of the debate, meaning that there won’t be a one-size-fits-all answer to the question of whether to harvest a voucher specimen. She suggests that when a biologist decides to collect a specimen or instead gather alternative evidence such as photographs and DNA samples, they should publish their reasoning behind that decision along with the species description. I present a (pseudo)formal framework for engaging in this sort of reasoning, presented as a utility function.

Opponents of routine harvest have typically argued that the scientific benefits of harvest are relatively small. But on its own, this fact doesn’t entail that the expected value of harvest is often negative. If the increased probability of extinction due to harvest is also relatively small, then the utility of harvest may be positive despite the scientific value of harvest being low compared to the disvalue of extinction.

But how are we to assess the increased probability of extinction? I examine two case studies, among the most prominent cases where taxonomists decided not to harvest a voucher specimen. I show how use of a technique from conservation biology known as Population Viability Analysis (PVA), conducted here using the freely available VORTEX individual-based modeling software, can be used to estimate k as a function of population size. The results of these two case studies, I argue, suggests: (1) that fears of voucher collection contributing to extinction are probably exaggerated, and (2) PVA provides a useful tool for researchers in both making and defending their decision on whether or not to lethally collect vouchers.

74. Evidence Hierarchies, Industry Bias and Dynamic Epistemic Environments
Daniel Saunders

Some evidential norms exhibit negative feedback effects which diminish their utility as they receive increasingly widespread uptake from the scientific community. This paper argues that the norm of privileging randomized controlled trials (RCTs) as the gold standard of medical evidence is one such example. One of the central tenants of the evidence-based medicine (EBM) movement is that RCTs belong at the top of the evidence hierarchy. While their privileged status has been criticized by a number of philosophers, less attention has been paid to the dynamic relationship between EBM's evidence hierarchy and the pharmaceutical industry. This paper explores both directions of the relationship - the effect that pharmaceutical funding has on EBM and the effect EBM has on pharmaceutical funding. I claim that the practice of privileging randomized controlled trials generates incentives for pharmaceutical companies to engage in research practices that ultimately undermine the original justification for the hierarchy.

Justifications for the gold standard status of RCTs typically appeal to their superior ability to rule out sources of confounding compared to observational studies. However, RCTs also receive disproportionate amounts of funding from the pharmaceutical industry. A wealth of evidence indicates that industry funding introduces multiple kinds of bias into the literature base, which skews the literature in favour of sponsors. These problems are particularly salient in the context of meta-analysis where industry funded RCTs tend to be overrepresented in the population of published literature. Privileging RCTs in the context of a research environment that receives large amounts of funding from industry can often lead to error.

The rise of EBM realigned research incentives for pharmaceutical companies. As the medical community increasingly viewed RCTs as the gold standard of evidence, gathering supportive RCTs became the most effective way for companies to market their drugs to medical professionals. They subsequently developed an institutional infrastructure for funding, conducting and publishing supportive RCTs in a way that aligned with the goals of marketing.

The medical research community can go a long way to solve these problems by placing a greater value on evidential diversity. Categorically privileging one type of study over others renders medical research fragile to manipulation. Evidential diversity considerations should include both the diversity of study designs but also the diversity of research communities who produce evidence. Placing greater emphasis on evidential diversity makes it more difficult for pharmaceutical companies to skew the literature in their favour and places a premium on seeking out evidence produced by independent researchers. This shift in values can easily be incorporated into meta-analyses. For example, researchers can perform subset analysis to see whether their conclusions are robust across multiple studies designs and sources of research funding.

75. Thinking with Ecosystems: Function, Dysfunction and the Notion of Disease

Tamar Schneider

In this talk, I will argue for the use of microbial ecology conceptual tools in the microbiome study in medicine and against the notion of dysbiosis. To argue that, I refer to the microbiome in the host as an ecosystem and its function as a causal role function.
The microbiome is an ecological community in its environmental niche, and in the body, it is connected to various pathological conditions, from inflammatory bowel diseases to obesity. This understanding of the microbial role in health and disease led to the accepted notion of the microbiome as an organ by itself (the gut microbiome in particularly). As such, the microbiome function in medical studies often referred to in terms of functional symbiosis or dysfunctional dysbiosis.

The concept of dysfunction presupposes definitions of a necessary function the trait/organ/organism has such that without it, they are dysfunctional. To have a necessary function for a trait, the function is connected to the reason the trait/organ/organism exists in the system. This framework of function in biology and physiology called the selected effect function (Wright 1976). The function is explained historically or etiologically and connected to the existence of the trait/organ/organism. Thus, a clear distinction of the identity of the entity/organ/ecosystem is needed (i.e., the identity is its function). In ecology, this definition is problematic. The definitions of dysfunctional ecosystems in ecology ascribe certain functions the ecosystem ought to have as a "normal' or within its normative capacities. However, attributing a dysfunction to an ecosystem narrows them to their desired functions instead of looking at their various interactions and interdependence with other ecosystems. Without such a definition, we can think of the ecosystem as possessing multiple functions within a web of interactions such that it reacts to changes and can change accordingly or die. Such openness and interactionist approaches shift the perspective of boundaries as well as the terminology of functions.

Similarly, in thinking about the microbiome, I argue that there can be more than one function and that functions are dynamic and can change. Thus, it is not possible to know all the necessary functions, only of some sufficient functions. Therefore, we cannot argue for dysfunction in general. We can argue for dysfunction in local and situated cases -- in a specific context with the desired function for a particular purpose that is not obtained. Acknowledging that this is a local claim, also has social and political values. For example, the microbial activity in the gut contributes to various systems such as the immune system and metabolism. It thus is not in a state of dysbiosis that defined sickness in conditions such as obesity or bowel syndromes. This perspective can also benefit other sections in medicine such as in women health. This, for example, prevents from looking at the uterus as dysfunctional after menopause (or for women who don't desire children), a belief that leads to higher rates of hysterectomy in benign conditions that have alternative treatments (Parker et al. 2009, 2013).

76. Bringing Different Voices into Science: Comparing Arguments from Feminist Philosophy and the Citizen Science Movement

Andrew Schroeder

Feminist philosophers of science have long argued for the importance of bringing feminist perspectives into scientific research. Proponents of citizen science argue for the importance of bringing the public into science. A key claim in both cases is that broadening participation in science will improve the quality of the resulting science. In this paper, I set aside what might be regarded as ancillary or indirect benefits (such as the social benefits that flow from having more women in science, or the increased public understanding of science that might be the result of citizen science) and focus narrowly on direct benefits to scientific research itself — respects in which scientific results may be improved through greater inclusion.

I argue that we can distinguish three respects in which both feminist philosophers and proponents of citizen science think science can be directly improved through greater inclusion: epistemic, ethical, and political. The *epistemic* improvement is the most familiar: greater inclusion can lead to scientific results that are more accurate or precise. (For example, feminist scientists may propose hypotheses that were not previously entertained by scientists in male-dominated fields; or the public may possess knowledge of local practices that enable more effective data collection.)

Once we recognize that science is value-laden, though, we can see that scientific results can be improved in two other ways. Results can be improved *ethically*, by being grounded in values that are ethically superior. (For example: feminist researchers may propose a definition of sexual assault that is based on an ethically superior understanding of consent.) I argue that there is also a distinct sense in which scientific results can be improved *politically*, by being grounded in values that are politically legitimate. (For example: by allowing the general public to set important scientific parameters, such as the social discount rate, the ensuing results can become a legitimate basis for policy-making, in a way that they otherwise would not have been.)

I conclude the paper by comparing case studies from the two literatures. I conclude that the epistemic and ethical reasons to bring feminist and citizen perspectives into science are similar in important respects. But I argue that the political reasons to bring the public into science are quite different from the political reasons to bring feminist voices into science. Incorporating the public and incorporating feminist voices can both increase the political legitimacy of the resulting science, but they do so in significantly different ways. I end by briefly considering what this tells us about other calls for greater inclusion in science — in particular, how we should respond to recent calls to include a wider range of political and religious viewpoints within science.

77. Medical Outliers, Intersectional Subjects, and the Peculiar Failure of the Differential: A Proposal for Improvement
Devora Shapiro

Zebras may be rare in comparison to horses, globally speaking, but in the right contexts (or continents) they are not so unlikely to be the source of hoofbeats. So too with our medical outliers – patients whose symptoms defy common explanation, and whose response to treatment fails to conform to standard expectations. In this paper I suggest that our difficulty in identifying and treating such patients – medical outliers – is rooted in a dissonance between the phenomena we hope to capture and the frameworks for representation and evaluation that we have available; on the one hand we have the statistically rare diseases or disorders that we wish to identify and treat, and on the other we have a framework for diagnostic methodologies that is based on likelihoods and wide-ranging commonalities.

Our diagnostic methodologies largely rely on large scale trends occurring in statistically significant population groups. Our rare diseases and disorders, however, are rare because they are statistically uncommon. In this presentation, therefore, I illustrate the need for alternative frameworks – structural representations - for the evaluation and diagnosis of uncommon diseases and/or medical conditions. Specifically, I discuss the benefits of approaching diagnosis and the diagnostic process through a reflexive model that considers the "constellation of symptoms" at one end, and on the other end conceptualizes patients as intersectional, multiply constructed subjects: biosocial phenotypes. In doing so, I demonstrate the potential for benefit that such an approach has to offer for the individual medical outlier. I further note the extended potential for such an approach to have application at the population level, particularly with regard to vulnerable, disadvantaged, or underserved unique and place-based communities.

I illustrate my point using the case of myalgic encephalomyelitis (ME) – a disease initially identified as "benign" and later as a psychosocial manifestation of hysteria, but that is currently understood to be a permanent, physically disabling condition related to inflammation of the brain and presenting with myriad verifiable and measurable symptoms. ME, having long been connected with Chronic Fatigue Syndrome (CFS), had until recently been repeatedly dismissed as a real or legitimate disease beyond the psychological. This dismissal has largely been attributed to three factors: 1) the predominance of the disease in females; 2) the variation in the presentation and epidemiology of the disease; and 3) the lack of explanation for the mechanism or etiology of the disease.

Having demonstrated the underlying issue as illustrated through the case of ME, I discuss the proposed framework and approach to diagnosis. The approach offered addresses the basic concern represented by statistically driven diagnostic algorithms and utilizes an achievable approach to information that can be implemented utilizing available technological resources. Further, I suggest it also offers flexibility with understanding "individual patients" as well as "patient populations," and promises improved responsiveness to concerns realized in vulnerable or disadvantaged communities and persons.

78. Plotting Syntheses of Rubber and of DNA: How Chemists use Chemical Structures and Reaction Pathways to Craft New Chemistries
Alok Srivastava

The abundant variety of robust DNA chemistries currently enabling the research worlds of Biotechnology and Synthetic Biology got their start in 1955 in the work of Michelson & Todd who demonstrated the laboratory synthesis of a dimer of DNA - a dinucleotide, two-unit polymer of DNA - through formation of a single inter-nucleotide linkage. This work demonstrated three inter-related scientific objects: (i) the specific chemical structure of the inter-nucleotide linkage from other alternative linkages, (ii) an actual chemical synthetic method - a recipe - for driving the preferential formation of this linkage from other possible reaction pathways, and (iii) the chemical pathway - the sequence of changes in chemical structure - for this preferential formation of the specific linkage between DNA nucleotides. In a few years this triad of chemical structure, chemical method and chemical pathway had been used, versioned, and elaborated in multiple laboratories so that some labs were making 20-mer polymers of well-defined sequences. The world of DNA chemistries was set into vigorous growth and development. Over the next 50 years, the world of DNA chemistries had grown to provide the DNA Synthesis economy of today's research and biomedical worlds.

This history of the growth and development of the chemist's world of DNA chemistries recapitulates the earlier history of the growth and development of Rubber chemistries and polymer chemistries. In 1884 the work of Tilden showed the dimerization for isoprene - an abundant chemical found in rubber sap - into dipentene. This work demonstrated (i) the chemical formula of the dimer dipentene, (ii) the chemical method of synthesis, and (iii) the chemical pathway of dimerization. This triad became the basis of many projects of chemists exploring the making of synthetic rubber. Over the next 50 years these synthetic rubber chemists developed the worlds of polymerization chemistry, rubber chemistry and chemistries of plastics.

In this paper I study the above two cases using the framework of narrative science being developed by Mary Morgan and colleagues and apply some principles from literary theory to explore the ways of world-making of chemists, i.e. how chemists use chemical structure and chemical methods and pathways to craft new chemistries. In this view, chemical structures function as narrative-plots and chemical methods (or recipes), and pathways, function as chronicles. Together they work in an interdependent and interactive manner like an engine for the elaboration of new 'versions and visions' of the world i.e. new chemistries.

Chemical structures denote and symbolize the arrangement of specific materials and locations of possible chemical change - i.e. the material order of chemical change. Chemical pathways denote and symbolize the sequencing of material and temporal change for specific classes of molecules. The specific chemical method is an example recipe, a chronicle and procedure of actual materials and contingent conditions that implement the particular chemical pathway. This is the complementary pair: narrative-plots and the conceptual and exemplary chronicles that together are an engine, like in the elaboration of literary genres, of a specific worlds of chemistries.

79. Assessing the Dynamics of Novel Tool Development: The Rodent Operant Touchscreen Chamber as a Case Study
Jackie Sullivan

In areas of neuroscience that investigate cognition, the development of innovative and reliable cognitive testing tools is equally as important as the development of tools for successfully intervening in, visualizing and decoding brain activity (e.g., Kraukauer et al. 2017). Yet, what kinds of normative constraints must cognitive testing tools satisfy in order to be successful? What measures do researchers take in order to satisfy these constraints? My aim in this talk is to address these questions via an analysis of one such state-of-the-art cognitive testing tool: the Bussey-Saksida Rodent Touchscreen Operant Chamber, which was first developed in the 1990s (Bussey, Muir, Robbins 1994; Bussey, Saksida, Rothblat 2001) and is now used in over 300 laboratories worldwide.

To provide some relevant background, computerized touchscreen tools for assessing cognitive impairments in humans with neurodegenerative or neuropsychiatric disorders have existed since the late 1980's (Sahakian & Owen 1992). While such tools serve important diagnostic and predictive functions, developing effective interventions for treating these impairments requires first testing such interventions in rodents. In order to do this, however, tools for assessing cognitive impairments in rodents needed to be developed. Conventional tools like the Morris Water Maze (spatial memory) and fear-conditioning paradigms were insufficient because they could not be used to assess the wide variety of cognitive functions shown to be impaired in human neurodegenerative and neuropsychiatric disorders. Tim Bussey, Lisa Saksida and their collaborators thus set out to develop a touchscreen apparatus that could be used to develop a flexible suite of rodent cognitive tasks that would mirror those computerized tasks already being used with humans. The aim of their research over the past 25 years has been to facilitate the translation of findings from animal models to humans, from the bench to bedside.

In this talk, I argue that the success of these tools for achieving these translational aims is contingent on them meeting a number of epistemic desiderata. For example, rodent tasks, like human tasks, must individuate discrete kinds of cognitive functions (construct validity) and ideally involve the same neural circuits in rodents and humans (neurocognitive validity). Ensuring that these requirements are met is an iterative process that includes both exploratory experiments (e.g., Feest 2012; Steinle 1997) to implement and test new tasks as well as hypothesis-driven experiments to test causal hypotheses in which task performance of wild-type (control) mice is compared to that of rodent models of neurodegenerative or neuropsychiatric disorders. Moreover, guaranteeing the reproducibility of the data across laboratories involves an unprecedented amount of coordination among members of the Bussey-Saksida research team as well as across the broader community of rodent touchscreen users (cf. Beraldo et al. 2019). I end by examining the nature of this coordination and addressing the question of whether it is a requirement for the success of other experimental tools in neuroscience—in particular interventionist tools.

Beraldo, F., Palmer, D., Memar, S., Wasserman, D., Lee, W. Liang, S., Creighton, S, Kolisnyk, B., Cowan, M., Mels, J., Masood, T., Fodor, C., Al-Onaizi, M., Bartha, R., Gee, T., Saksida, L., Bussey, T., Strother, S., Prado, V., Winters, B., Prado, M. 2019. MouseBytes, an Open-access High-throughput Pipeline and Database for Rodent Touchscreen-based Cognitive Assessment. *Elife* 8.

Bussey, T., Muir, J., and Robbins, T. 1994. A Novel Automated Touchscreen Procedure for Assessing Learning in the Rat using Computer Graphic Stimuli. *Neuroscience Research Communications* 15(2).

Bussey, T., Saksida, L. and Rothblat, L. 2001. Brief Communications-Discrimination of Computer-Graphic Stimuli by Mice: A Method for the Behavioral Characterization of transgenic and Gene-Knockout Models. *Behavioral Neuroscience* 115(4).

Feest, U. 2012. Exploratory Experiments, Concept Formation and Theory Construction in Psychology. In *Scientific Concepts and Investigative Practice*, Springer, pp. 167-189.

Krakauer, J., Ghazanfar, A., Gomez-Marin, A., Poeppel, D. 2017. Neuroscience Needs Behavior: Correcting a Reductionist Bias. *Neuron* 93(3).

Sahakian, B. and Owen, A., 1992. Computerized Assessment in Neuropsychiatry using CANTAB: Discussion Paper. *Journal of the Royal Society of Medicine* 85(7).

Steinle, F. 1997. Entering New Fields: Exploratory uses of Experimentation. *Philosophy of Science* 64.

80. The Demise of in Vivo Experimentation and the Problem of Methodology Elimination in Toxicological Risk Assessment
Richard Sung

In 2019, the US EPA announced that all funding for in vivo experimentation in toxicological risk assessment will cease by 2035. The rationale is that in vivo experimentation is costly, time-consuming, unreliable, and ethically contested. Nevertheless, the EPA's decision has led to a heated controversy concerning the epistemic and ethical status of in vivo experimentation, the reliability of emerging methodological alternatives, and the nature of toxicological risk assessment as both scientific and public policy practice.

In this talk, I situate the controversy within the ongoing philosophical debate on the roles of values in science. In particular, I critically engage with Kevin C. Elliott's recent discussion of the future of the adverse outcome pathway (AOP) framework in toxicological risk assessment. The AOP framework, developed in response to the 2007 National Research Council (NRC) report Toxicity Testing in the 21st Century, is intended to organize available biological information along various levels of biological organization and illuminate causal relationships between processes at the molecular level and processes at cellular, organ, organ-systemic, and organism levels. The AOP framework has been endorsed as a major alternative to animal models due to its potential to provide regulators the causal understanding necessary for predicting the potential toxicity of chemicals without demanding expensive, time-consuming, and controversial animal studies.

While recognizing the AOPs' revolutionary potential, Elliott argues for a serious consideration of their development and deployment in the context of intense social and political polarizations surrounding environmental regulation. In particular, Elliott emphasizes that any assumptions, gaps, or uncertainties concerning AOPs may generate contestations that could compromise the AOPs' potential, resulting in detrimental consequences for relevant stakeholder groups. Recognizing such possible "gridlocks", Elliott proposes two arguments concerning the future of AOPs. First, the quality and social utility of AOPs can be maximized if their development processes remain true to two principles: transparency and participation. Second, AOPs are likely to be most fruitful in "win-win" contexts where relevant stakeholders have no incentives to emphasize uncertainties and limitations of AOPs.

By combining my own study of the AOP framework's socio-political context with insights drawn from philosophical discussions on the legitimate roles of values in science, I argue as follows. First, while the argument from inductive risk and the history of toxicological risk assessment do highlight the necessity of considering socio-political contexts in which AOP frameworks are situated, transparency and participation do not suffice as the guiding principles for future AOP development and deployment when there is the possibility of the corruption of scientific practices. This is especially so when the socio-political context one faces can be characterized as an instance of the Ibsen predicament elaborated by Daniel Steel. Second, despite their initial appeal, "win-win" contexts are vulnerable to two scenarios: 1) when the relevant stakeholders fail to establish and maintain a common ground for prioritizing particular values over others and 2) when the establishment of agreement among the stakeholders necessitates epistemic accommodation, namely pursuing agreement at the cost of addressing significant epistemic limitations of a research methodology. Based on the two arguments, I argue that the EPA's decision to eliminate in vivo experimentation is problematic in that it could generate supposedly “win-win” contexts that involve epistemic accommodation rather than meaningful deliberation on research methodologies constituting risk assessment.

81. Refugee Mental Health and Psychotherapy Chatbots: Learning from ELIZA
Şerife Tekin

Gulf between the needs of individuals with mental disorders and available mental health care services disproportionately affects vulnerable populations with a high risk of developing mental health problems, such as refugees. Advances in applications of artificial intelligence and the use of data analytics technology in biomedicine create some optimism, as these technologies may fill the need-availability gap by increasing mental health care resources for refugees. One resource is smartphone psychotherapy chatbots, i.e., artificially intelligent bots that offer cognitive behavior therapy to their users with the aim of helping them improve their mental health. Proponents of using smartphone psychotherapy chatbots as therapists often list their low cost, wide accessibility through cell phones, and availability in different languages as advantages and argue that these make them an ideal medical tool, especially in areas with a shortage of therapists who speak the native language of refugees requiring care.

While a number of studies have highlighted the positive outcomes of using smartphone psychotherapy chatbots to handle anxiety related problems, no conclusive data illustrate their effectiveness or warrant their use for refugee mental health care. In addition, while some philosophers of science and bioethicists started examining various philosophical and ethical issues surrounding psychotherapy chatbots, they have not yet benefited from the history of Artificial Intelligence research to inform conversations about effectiveness of psychotherapy chatbots. In this paper, taking cues from Joseph Weizenbaum’s own views on ELIZA –the computer program created to demonstrate the superficiality of communication between humans and machines—I will raise epistemic and ethical concerns about the efficacy of treating refugees with psychotherapy chatbots. I will focus on the specific features of smartphone psychotherapy chatbots designed to address refugee mental health (such as an Arabic speaking bot Karim) to determine whether they have genuine promise for ethically addressing mental health needs of refugees. Despite their attractiveness they contain significant risks, even obstacles, to effective and ethical treatments. In addition to epistemic and ethical constraints of bots in addressing mental health challenges, I worry that motivating the development of this technology to address the growing needs of refugee populations medicalizes social and political problems. It encourages masking these instead of proposing solutions.

82. Value Integration in Climate Impact Modelling: Towards a Transdisciplinary Approach
Henrik Thorén

In assessing the wider impacts of climate change on social and economic systems scientists have in recent years increasingly relied on the use of climate impact models such as integrated assessment models (IAM). Studies based on such models provide the basis for much of the results synthesised by the third working group of the IPCC and are furthermore widely used as a basis for emission scenarios that feed into climate projections.

The models themselves build on the integration, in a range of different ways, of various sub-models (or modules) in order to capture the dynamics of larger systems. Many IAMs couple a comparatively simple climate model to a macro economic model but the models can be complex consisting of hundreds of equations describing a range of different systems and economic sectors.

The models and their use has however been deeply controversial, even among economists, and some have even argued that they are largely useless and that (neoclassical) economists should be banned from the IPCC altogether. The approach, critics would have it, wildly underestimates the costs of climate change and as well as the benefits of mitigation, obscure a host of relevant uncertainties, and commit to idealisations and simplifications that severely deteriorate the reliability of outputs. A common contention is that the use of these models involve serious moral hazards as they appear to justify mitigation policies that are woefully inadequate.

How are these models best put to use? And what can they tell us about how to best craft policies in the context of climate change? In this paper we analyse these questions in terms of value integration; and in particular how (relevant) none-epistemic values are to be appropriately treated.

The aim in this paper is to, by drawing on recent literature on the role of values in science (e.g. Elliott 2017; Douglas 2009; Parker and Winsberg 2018; Intemann 2015), analyse what role climate impact models could and should play in policy making and discuss some of the limitations with these models. We argued that the emphasis has for the last few decades been mostly on technical aspects of the integration of sub-models and modules with one another, rather than developing a robust transdisciplinary process—that is to say, a practice that involves relevant stakeholders and constituents—that can take into account the appropriate non-epistemic (contextual) values and resolve trade-offs between epistemic and non-epistemic values. We suggest that this is an important reason why it remains largely unclear how the model outputs can or should inform policy making efforts.

In the last part of the paper we outline, based on the recent literature on transdisciplinarity, what such a practice might look like and what the implications of deploying it would be for the model development process. In particular, the increasing complexity of the models and the continual interdependency between model structures and how they are used, appear to provide an important constraint.

83. Uniqueness and Biological Individuality
Rose Trappes

Are all biological individuals unique? This question has been treated in metaphysics for centuries, and in philosophy of biology for several decades (Strawson 1959; Santelices 1999; Lowe 2003; De Sousa 2005; Clarke 2010). When it comes to qualitative properties, philosophers of biology and metaphysicians tend to answer with a fairly certain "no." Biological individuals often have identical genomes, and many seem to be exactly phenotypically identical too. And even if we don't have real examples of qualitatively identical biological individuals, it is conceivable that they could exist. Relational uniqueness, such as having a unique spatiotemporal position, seems more likely to be necessary for biological individuals, but even this could be and is debated.

In contrast with philosophers, biologists believe that biological individuals are both qualitatively and relationally unique. In this talk I present data from my qualitative research on an interdisciplinary group of biologists studying animal behaviour and individualised niches. In questionnaires and interviews, biologists regularly associate individuality with uniqueness and frequently assert that animals and other organisms have unique phenotypes, experiences, and niches. What should we make of this? Are biologists just mistaken? Or is uniqueness really necessary for biological individuality? To answer this question I look at biologists' justificatory practices and the role their beliefs play in their methodological choices.

The biologists justify their beliefs in uniqueness in two ways. First, they appeal to evidence from everyday experience with animals in the lab and field as well empirical research showing that phenotypic differences between individuals arise despite identical genotypes and near-identical environmental conditions (Freund et al. 2013; Bierbach, Laskowski, and Wolf 2017). Second, biologists use theoretical arguments to reason that exactly identical biological individuals couldn't exist because the causal processes behind phenotypes, experience and niches are too complex and sensitive to be exactly duplicated. Such theoretical reasoning, together with empirical evidence, provides strong support for the biologists' claims that all biological individuals are unique. Specifically, I argue that biologists are claiming that uniqueness is biologically necessary for biological individuals, regardless of any metaphysical possibility of non-unique individuals. This claim is well supported, and I conclude that philosophers too should accept that all biological individuals are unique, with the caveat that this claim carries biological but not metaphysical necessity.

There remains a worry that the beliefs biologists express about uniqueness are largely epiphenomenal to their work, a consequence of conversing with a philosopher rather than engaging in scientific practice. I therefore close by suggesting that beliefs about uniqueness serve as a guide for biologists to adopt methodologies that look closer at individuals and the differences between them. Such methodologies can range from merely studying groups below the population level or examining variation around a mean in a standard experiment, through to individual-based modelling and linear regression. These disparate models share the intention to get closer to the individual, guided by the belief that all individuals are unique and that some aspects of this uniqueness should be studied as it may be biologically significant.

Bierbach, D, K. L. Laskowski, and M. Wolf. 2017. Behavioural Individuality in Clonal Fish Arises despite Near-Identical Rearing Conditions. *Nature Communications* 8(1).

Clarke, E. 2010. The Problem of Biological Individuality. *Biological Theory* 5(4).

De Sousa, R. 2005. Biological Individuality. *Croatian Journal of Philosophy* 5(14).

Freund, J., A. M. Brandmaier, L. Lewejohann, I. Kirste, M. Kritzler, A. Kruger, N. Sachser, U. Lindenberger, and G. Kempermann. 2013. Emergence of Individuality in Genetically Identical Mice. *Science* 340 (6133).

Lowe, E. J. 2003. Individuation. In *The Oxford Handbook of Metaphysics*, edited by M. J. Loux and Zimmerman, D. W., 75–95. Oxford University Press.

Santelices, B. 1999. How Many Kinds of Individual Are There? *Trends in Ecology & Evolution* 14 (4).

Strawson, P.F. 1959. *Individuals: An Essay in Descriptive Metaphysics*. Routledge.

84. Using Metaphysics to Institutionalize Pluralism
Katherine Valde

There is a growing recognition among scientists and philosophers that metaphysical presuppositions guide scientific research. These metaphysical presuppositions guide scientific research by articulating an ontology for a particular domain of phenomena, that is, by making a claim about what sort of things there are and what they are fundamentally like. These ontological claims, in turn, prescribe a particular methodology for how to go about investigating and explaining those kinds of things. There is thus what I call a move from metaphysics to methods.

A key question that has yet to be addressed is what sort of attitude ought to be taken towards such metaphysical presuppositions. Scientific research cannot take place in a metaphysical vacuum, and there are a variety of different attitudes one can take towards these frameworks, as exemplified by recent cancer research.

This presentation will begin by examining how current cancer research, which is largely centered on Somatic Mutation Theory (SMT), grows out of a mechanistic metaphysical picture of the world and how that mechanistic framework limits methodological choices to those that the make use of mechanistic concepts. Next, I will examine the newly proposed processual alternative to this approach. The processualists argue that mechanistic accounts are inadequate because they presuppose an incorrect metaphysical picture of the biological world. I will argue that there is an important problem with the argument from metaphysics to scientific methods. For example, the processualists are calling for an unwarranted limitation of our scientific methods. My objection is not to processualist research per se, but to grounding cancer research exclusively on a any singular metaphysical framework. Instead, I argue that metaphysical frameworks are underdetermined by empirical research, and thus we should turn to the important pragmatic aims of science to ground a genuine explanatory and methodological pluralism for cancer research.

While an agnostic attitude towards metaphysical frameworks could help to ground a genuine pluralism of research frameworks placing this metaphysical burden on individual researchers misses the more systemic issues causing siloed research agendas. This presentation will argue that interventions at a higher level will be more effective in bringing about pluralism in practice. For example, providing training on the underdetermination of metaphysical frameworks by empirical research and inductive risk to grant reviewers could help to ensure a diversity of programs operating on a variety of metaphysical frameworks are funded and executed.

I close my argument by highlighting the importance of institutionalizing the pluralism. Given the inductive risk for large non-epistemic consequences in cancer research, pragmatic values need to play a role in our reasoning about what sorts of research to pursue. Rather than making these metaphysical arguments to individual researchers, making them to government and industry sources of funding is the most promising route to achieving a productive pluralism.

85. Constitutive Relevance in Cognitive Science: The Case of Eye Movements and Cognitive Mechanisms

Dingmar van Eck

In this contribution I assess whether the "No De-Coupling" (NDC) theory of constitutive relevance in mechanisms (Baumgartner and Casini 2017) is a useful tool to reconstruct constitutive relevance investigations in scientific practice. The NDC theory has been advanced as a framework theoretically superior to the mutual manipulability (MM) account of constitutive relevance in mechanisms (Craver 2007) but, in contrast to the MM account, has not yet been applied to detailed case studies. I apply the NDC account to two case studies drawn from cognitive science research on the role of eye movements in mechanisms for cognitive capacities: i) pattern copying (Ballard et al. 1997) and ii) visuospatial memory retrieval (Johansson and Johansson 2014). I argue that the NDC account is also applicable to empirical practice and that it fares better than the MM account on both theoretical and empirical grounds.

On the MM account, in brief, the activity of an entity is constitutively relevant for a macro-level behavior or capacity if one can change the macro-level behavior by intervening to change the entity's activity, and if one can change the activity of the entity by intervening to change the macro-level behavior. So, when macro and micro levels are mutually manipulable, i.e., difference makers for one another, the inference is warranted that the micro level is constitutively relevant for the macro level. A major alleged virtue of the MM account is that it gives a faithful reconstruction of top-down and bottom-up experimental practices in the cognitive and neural sciences.

Recently, however, the MM account has come under heavy fire: given the assumption of constitutive dependencies between micro and macro levels, (ideal) interventions that induce changes in one level by inducing changes in the other are impossible in principle. Micro and macro levels, in case of constitution, can only be manipulated via fat-handed interventions that cause changes at both levels via separate causal paths.

The NDC theory of constitutive relevance pivots on this insight. According to NDC, in a nutshell, evidence for constitutive relevance hinges on the notion that, in case of constitution, changes in micro and macro levels can only be induced through the same (fat handed) intervention along separate causal paths, whilst interventions that would break this coupling, i.e., ones that only induce changes at the macro level, and not in one (or more) constituents at the micro level, are impossible. When extensive testing suggests that the latter interventions are not possible, this is taken to provide abductive evidence for constituency relations between the micro and macro level.

Although NDC is a major theoretical improvement on the MM account, it is still an open question whether NDC accords with and can be used to elucidate constitutive relevance assessments in scientific practice. Up until now, it has only been illustrated in terms of toy examples. In this contribution I address this question by applying the NDC account to the two above-mentioned case studies.

Ballard ,D.H., Hayhoe, M.M., Pook, P.K., & Rao, R.P.N. 1997. Deictic Codes for the Embodiment of Cognition. *Behavioral and Brain Sciences* 20.

Baumgartner, M., and Casini, L. 2017. An Abductive Theory of Constitution. *Philosophy of Science* 84(2).
Craver, C.F. 2007. *Explaining the Brain: Mechanisms and the Mosaic Unity of Neuroscience.* Oxford University Press.
Johansson, R., and Johansson, M. 2014. Look Here, Eye Movements play a Functional Role in Memory Retrieval. *Psychological Science* 25(1).

86. Beyond Actual Difference Making: Toward a Better Interventionist Account of Causal Selection
Zina Ward

Not all causes are created equal. This is the idea behind the burgeoning literature on causal selection, which aims to identify features that make some causes more important than others. In his classic paper on causal selection, Waters (2007) argues that biologists implicitly make a distinction between the causes of a trait and the subset of those causes that bring about "actual differences" in the trait. He calls the latter "actual difference makers" (ADMs) and proposes an interventionist characterization of them (Woodward 2003). Waters argues that biologists privilege ADMs and are right to do so. What's more, he claims, genes are the most causally specific ADMs in biology and are therefore ontologically privileged over other causes.

Although Waters' defense of gene-centric biology is highly controversial (Stegmann 2012, Griffiths and Stotz 2013), the concept of an ADM has not itself been challenged. The prevailing opinion is that, whatever the problems with Waters' particular application of actual difference making, the concept itself is a useful tool for causal selection. My aim in this paper is to argue that this is mistaken. The problems with actual difference making run deep.

First, I show that Waters' definitions of "the" and "an" ADM are subject to several ambiguities and infelicities. For instance, Waters does not specify whether actual difference making is meant to be a type- or token-causal notion. This matters because type- and token-causal claims are analyzed differently within interventionism. Next, I suggest that the pre-existing concept that Waters is attempting to analyze is that of a component of variance. Components of variance are identified through analysis of variance (ANOVA) methods, which partition observed variability and attribute it to specific factors. Interpreting ADMs as components of variance helps resolve some of the ambiguities previously discussed but also brings to light two additional problems: it undermines condition (iv) in Waters' definition of an ADM, and it implies that the concept of actual difference making can only be applied when causes contribute additively to the effect (Lewontin 1974).

Finally, I offer a diagnosis: many of the above problems arise because the phrase "partially account for" invokes a "hydraulic" model of causation that is fundamentally incompatible with Waters' commitment to interventionism (Burge 2007). When we say that variation in one variable accounts for variation in another, we are implicitly appealing to a picture of causation on which there is a certain amount of "energy" required to "get the [causal] job done" (ibid., 380). To identify causes that account for variation is to determine different sources of the finite "causal liquid" responsible for the variation. Under interventionism, however, the number of causes that are potentially relevant to an event is in principle unlimited. Since there is no interventionist analogue of the metaphor of a finite pool of causal liquid, the phrase “partially account for” resists interventionist analysis.

Because of this tension, I argue that Waters’ concept of actual difference making should be rejected. The tools one uses for causal selection must be compatible with the general theory of causation that one adopts.

87. Kinship Unbound: Constructivism, Bioessentialism, and the Ethnographic Record
Rob Wilson

In recent years, the study of kinship has made a reinvigorated return within anthropology, both through innovative work in the "new kinship studies" (Bamford 2019, Sahlins 2013, Strathern in press) that adopts a constructivist or performativist view of kinship, and in research that is closer in spirit to the bioessentialist tradition in kinship studies (McConvell et al. 2013, Shapiro 2018, Wilson 2016). In this talk I want to critically explore a key argument for constructivism about kinship that appeals to the cross-cultural diversity one finds in the practice of kinship. The basic idea of the argument is that, given the variety of non-biological ways in which people practice or perform kinship across different cultures, bioessentialism about kinship is much less plausible than is constructivism. After briefly characterizing the shift from bioessentialism to constructivism (or performativism) in kinship studies (section 2), I state more fully the ethnographic argument for performativism that I will focus on and that has continuing influence in the study of kinship within cultural anthropology (section 3). The remainder of the talk critiques each of the three premises in this argument, initially by identifying problems with the claim that culture is conceptually prior to biology vis-à-vis kinship (section 4), which has been stated in several ways by Marshall Sahlins in his influential recent book *What Kinship Is*. I shall then provide reasons for rejecting the claim that cultural ignorance about biology favors performativism about kinship (section 5). Finally, I will challenge the pluralistic strand to performativism by drawing on the recent revival of extensionism about kinship terminologies (section 6). Since the views and claims challenged here have widespread, uncritical acceptance within cultural anthropology, this critique should have interest beyond the acceptance or rejection of performativism about kinship. It raises questions about methodology, about the relationship between theory and evidence, and about the prospects for an integrated view of kinship that moves beyond the present divide between constructivists and bioessentialists.

Bamford, S. ed. 2019. *The Cambridge Handbook of Kinship*. Cambridge University Press.
McConvell, P., I. Keen, and R. Hendery. eds. 2013. *Kinship Systems: Change and Reconstruction*. University of Utah Press.
Sahlins, M. 2013. *What Kinship Is--And Is Not*. University of Chicago Press.
Shapiro, W. 2018. ed. *Focus and Extension and the Study of Kinship: Essays in Memory of Harold W. Scheffler*. Australian National University Press.
Strathern, M. in press. *Relations: An Anthropological Account*. Duke University Press.
Wilson, Robert A. 2016. Kinship Past, Kinship Present: Bio-Essentialism and the Study of Kinship. *American Anthropologist* 118(3).

88. How Scientific Generalizations Facilitate Investigations

Yoshinari Yoshida

What value do generalizations have in science? Traditional philosophy of science would answer this question by appealing to contributions of generalizations to scientific explanations. A generalization is valuable because it enables us to explain phenomena, e.g., by providing a crucial premise for a deductive reasoning, serving as a basis under which different phenomena are subsumed, or indicating invariance of causal relationships (Hempel, 1966; Kitcher, 1981; Woodward, 2001). Another traditional approach is to treat generality as a basic virtue in science as some authors do when discussing epistemic/cognitive values or modeling desiderata (e.g., Kuhn, 1977).

Although the contribution to explanation and the intrinsic value of generality account partially for why generalizations are sought in science, they do not exhaust the reason why generalizations are valuable. In this paper, I discuss several ways generalizations facilitate investigations in research communities. Established or supposed generalizations describe interesting distributions of certain properties that are to be explained; they also serve as a basis for extrapolation across different systems by articulating similarity between them. Furthermore, I argue that search for generalizations is often beneficial even if it ends up in a failure to establish unified generalizations because it can promote communication and mutual informing among different subcommunities, which in turn promotes individual studies.

As an example, I examine recent progresses in the study of collective cell migration in developmental biology. Molecular and cellular aspects of collective migration have been elucidated through studies of various, seemingly unrelated systems of different taxa, such as fruit fly border cells, zebrafish lateral line, vertebrate vasculature, and invasive and metastatic cancer (e.g., Scarpa and Mayor, 2016). Researchers have compared collective migration in different systems and asked what common mechanisms and principles can be identified. This has identified cellular and molecular interactions shared across some systems, whereas it has also revealed interesting heterogeneity. Due to such heterogeneity, substantive, overarching models of collective migration have not been established. However, the framework of collective cell migration and comparison of various systems remain to be popular today. I discuss how this search for general mechanisms and principles has facilitated investigations in the field. It has articulated similarity between certain biological systems and thereby provided resources and bases for further research. Furthermore, the search for common aspects across the diverse systems has helped to connect otherwise unrelated subdomains of knowledge and promotes mutual informing among them, which has led to better understanding of each migration mechanism and more efficient investigations of the individual systems.

Hempel C. 1966. *Philosophy of Natural science*. Prentice Hall.
Kitcher P. 1981. Explanatory Unification. *Philosophy of Science* 48(4).
Kuhn T. 1977. Objectivity, Value Judgment, and Theory Choice. In: *The Essential Tension: Selected Studies in Scientific Tradition and Change*, The University of Chicago Press, pp 320–339.
Scarpa E., Mayor R. 2016. Collective cell migration in development. *Journal of Cell Biology* 212(2).

Woodward J. 2001. Law and Explanation in Biology: Invariance is the Kind of Stability that Matters. *Philosophy of Science* 68(1).

89. Cancer Research Suggests Extending the Concept of Modeling
Martin Zach

The philosophy of scientific modeling is now a well-established topic. After abandoning the syntactic and semantic conceptions in favor of the practice approach (Morrison and Morgan (1999); Bailer-Jones (2009)), a widely accepted characterization of models has emerged. According to this view, “the modeler’s first move is the specification and investigation of a hypothetical system, or structure,” while the second move pertains to “consideration of resemblance relations between this hypothetical system and the real world (Godfrey-Smith (2006), p. 730; see also Weisberg (2007), p. 209 for essentially the same view). In this sense models have been construed as means of indirectly representing their target systems (Mäki (2009); Knuuttila (2011); Salis (2016); Frigg and Nguyen (2017)). Indeed, the practice of indirect representation is purportedly what distinguishes modeling from other ways of theorizing such as “abstract direct representation” (Godfrey-Smith (2006); Weisberg (2007)).

It is argued here that the received view sketched above does not adequately capture numerous modeling practices which can be collectively referred to either as “data-derived models” or “models derived from experiments.” The dominance of the received view is also quite puzzling, given that there is an already existing vast literature on, e.g., Big Data modeling, which doesn’t seem to fit the framework described above. Or consider models in experimental structural biology (Mitchell and Gronenborn (2017)). Consequently, it is not the case that models are necessarily something that is being investigated instead of their target systems as many have claimed. Nor is it the case that all models are constructed on the basis of considering certain modeling assumptions. Finally, it is not the case that all models are indirect representations.

These claims are defended by describing the multiple layers of an experimental practice that stands behind deriving mechanistic models in cancer research. More specifically, I will discuss the ways in which the role of myeloid-derived suppressor cells (MDSCs) in cancer dissemination is being modeled. MDSCs are immature cells of myeloid origin that are associated with immunosuppressive effects (Gabrilovich and Nagaraj (2009)). Their origin, metabolic programme as well as their role in the tumor microenvironment has been the subject of cancer research in the past decade (Ostrand-Rosenberg and Sinha (2009); Kumar et al. (2016); Veglia et al. (2018)). Beyond their presence in the tumor microenvironment they have also been implied in facilitating cancer dissemination (Wang et al. (2019)).

I will describe a set of ongoing experiments investigating various aspects of MDSCs which aim at constructing a model of cancer dissemination that can account for and explain the data within a broader theoretical framework. These experiments provide a mosaic picture and include experiments on cell cultures (cancer cell lines), immunohistochemistry, animal tumor models and imaging techniques among other things. Many of these are neither completely hypothesis-free, nor fully explorative. Importantly, such model construction is far removed from what the received view would have us believed since the model is not built by considering certain assumptions but rather derived from the experimental results.

Furthermore, these practices can readily be distinguished from what has been called “data models”, and, using the terminology of Bogen and Woodward (1988), can legitimately remain to be identified with the construction of models of phenomena.

Bailer-Jones, D. 2009. *Scientific Models in Philosophy of Science*. University of Pittsburgh Press.
Bogen, J. and Woodward, J. 1988. Saving the Phenomena. *The Philosophical Review*, 97.
Frigg, R. and Nguyen, J. 2017. Models and Representation. In L. Magnani and T. Bertolotti (eds), *Springer Handbook of Model-Based Science*. Springer, pp. 49–102.
Gabrilovich, D. I. and Nagaraj, S. 2009. Myeloid-Derived Suppressor Cells as Regulators of the Immune System. *Nature Reviews Immunology*, 9.
Godfrey-Smith, P. 2006. The Strategy of Model-Based Science. *Biology & Philosophy*, 21.
Knuuttila, T. 2011. Modelling and Representing: An Artefactual Approach to Model-Based Representation. *Studies in History and Philosophy of Science Part A*, 42.
Kumar, V., Patel, S., Tcyganov, E. and Gabrilovich, D. I. 2016. The Nature of Myeloid-Derived Suppressor Cells in the Tumor Microenvironment. *Trends in Immunology*, 37.
Mäki, U. 2009. MISSing the World. Models as Isolations and Credible Surrogate Systems. *Erkenntnis*, 70.
Mitchell, S. D. and Gronenborn, A. M. 2017. After Fifty Years, Why Are Protein X-Ray Crystallographers Still in Business? *The British Journal for the Philosophy of Scienc*e 68.
Morrison, M. and Morgan, M. S. 1999. Models as Mediating Instruments. In M. S. Morgan and M. Morrison (eds), *Models as Mediators*. Cambridge University Press, pp. 10–37.
Ostrand-Rosenberg, S. and Sinha, P. 2009. Myeloid-Derived Suppressor Cells: Linking Inflammation and Cancer. *Journal of Immunology*, 182.
Salis, F. 2016. The Nature of Model-World Comparisons. *The Monist*, 99.
Veglia, F., Perego, M. and Gabrilovich, D. 2018. Myeloid-Derived Suppressor Cells Coming of Age. *Nature Immunology*, 19.
Wang, Y., Ding, Y., Guo, N. and Wang, S. 2019. MDSCs: Key Criminals of Tumor Pre-Metastatic Niche Formation. *Frontiers in Immunology*, 10.
Weisberg, M. 2007. Who Is a Modeler? *The British Journal for the Philosophy of Science* 58.