The Magazine

June 2023

Introduction

Hello PiBM network members!

If you can believe it, we are now halfway through the year, quickly entering the summer (or winter) and, for many, a much-needed break. Whether or not you're relaxing somewhere in the world, you'll find much to read in this month's newsletter. Apart from even more postgraduate possibilities to keep your grad students working on applications, there are various accouncements and fascinating events in the PiBM world to report. We even have not one but two interviews!

Enjoy!

Post-graduate Opportunities

PhD & Postdoc in Philosophy of Life Sciences – Bielefeld

PhD position (3 years) in the Philosophy of the Life Sciences at Bielefeld University, Germany. The Faculty of History, Philosophy und Theology, Department of Philosophy (research group philosophy of science, Prof. Dr. Marie I. Kaiser/Prof. Dr. Lara Keuck/Jun.-Prof. Dr. Alkistis Elliott-Graves).

Your tasks:

- conduct independent research in the philosophy of the life sciences and in related areas (65 %)
- teaching at Bachelor and/or Master level in philosophy and/or in the Master's program "Interdisciplinary Studies of Science"; 2-3 courses per year, in German or English (25 %) (cont...)

Post-graduate Opportunities (cont.)

- actively participate in the meetings and events of the philosophy of science group and the Department of Philosophy (5 %)
- organizational tasks that are part of the self-administration of the university (5 %).

For more information about our expectations, documents required for application and the online application form please have a look at the following website, English: https://uni-bielefeld.hr4you.org/job/view/2528/research-position-in-the-philosophy-of-the-life-sciences?page_lang=en and German: https://uni-bielefeld.hr4you.org/job/view/2530/wissenschaftliche-r-mitarbeiter-in-m-w-d-in-der-philosophie-der-lebenswissenschaften?page_lang=de

Application deadline: July 5th, 2023
Contact: Prof. Dr. Marie I. Kaiser
kaiser.m@uni-bielefeld.de

Postdoc position (3 years) in the Philosophy of the Life Sciences at Bielefeld University, Germany. The Faculty of History, Philosophy und Theology, Department of Philosophy (research group philosophy of science, Prof. Dr. Marie I. Kaiser/Prof. Dr. Lara Keuck/Jun.-Prof. Dr. Alkistis Elliott-Graves).

Your tasks:

- conduct independent research in the philosophy of the life sciences and in related areas (50 %)
- teaching at Bachelor and/or Master level in philosophy and/or in the Master's program "Interdisciplinary Studies of Science"; 4 courses per year, in German or English (25 %)
- contribute to writing grant applications and to organizing workshops/ conferences (15 %)
- actively participate in the meetings and events of the philosophy of science group and the Department of Philosophy (5 %)
- organizational tasks that are part of the self-administration of the university (5 %).

(cont...)

Post-graduate Opportunities (cont.)

For more information about our expectations, documents required for application and the online application form please have a look at the following website: https://uni-bielefeld.hr4you.org/job/view/2526/research-position-postdoc-in-the-philosophy-of-the-life-sciences?page_lang=en

Application deadline: July 13th, 2023
Contact: Prof. Dr. Marie I. Kaiser
kaiser.m@uni-bielefeld.de

Bordeaux & Marie Curie Postdocs

The Conceptual Biology & Medicine team at ImmunoConcept in Bordeaux would like to announce that it is still interested in hosting postdoctoral candidates through the Marie Sklodowska-Curie fellowship scheme, which is now accepting calls until September 13.

For more details about the fellowships, please see: https://marie-sklodowska-curie-actions.ec.europa.eu/actions/postdoctoral-fellowships For more details about the team at Bordeaux, either visit their website (https://immunoconcept.cnrs.fr/conceptual-biology-medicine/), or contact anyone on the team with whom you would like to work or further discuss this post-doc opportunity: Maël Lemoine, Thomas Pradeu, Fridolin Gross, Jan-Pieter Konsman, or Jonathan Sholl.

Jobs

Tenure-track job at University of Sydney

University of Sydney is offering some tenure-track jobs that are mostly research to begin with through the Sydney Horizon Fellowships. The focus should be health, climate and/or sustainability, but they are supposed to be open to all subjects. For those interested, the History of Philosophy and Science Department is a possible host, so please circulate among any underemployed early career candidates (less than 10 years post-PhD). Please note that the deadline is extremely close (July 5)!

Upcoming Conferences & Workshops

Egenis Workshops: Science in Public Early Career Workshop

Due to the global pandemic, Science in Public has been unable to run a full conference since 2019, but is now being re-launched! This workshop creates the first opportunity for new generations of 'early-career' researchers and professionals to come together: with each other, senior colleagues, SiP Committee members, and expert practitioners, to discuss the present and futures of our community in the UK. Workshop will be on July 14 from 11:00-17:00 BST, at the Peter Chalk Centre, University of Exeter.

Sessions include:

- Plenary Panel: Intersectional & Inclusive Approaches to Science Engagement (Emily Dawson, University College London; Hana Ayoob, Science Communicator and Illustrator; Ernesto Schwartz-Marin, University of Exeter; Tom Aechtner, University of Queensland; Stephen Jones, University of Birmingham)
- Roundtable: Pathways & Roadblocks to Connecting Research & Practice
- Discussion From the Floor: (Re)Shaping the Future of Science in Public
- Networking Lunch and Evening Social

See the Science in Public website https://scienceinpublic.org/ for more details and a full program. For more details and information about these Egenis workshops, see this page: https://sociology.exeter.ac.uk/research/sts/egenis/activities/events/

Announcements

Two ANR (French National Research Agency) projects awarded, Bordeaux

Maël Lemoine will be leading MEASURAGE, an ANR project on the theoretical framework necessary to correctly measure the progression of aging. Partners are Adrien Barton (Université de Toulouse) and Alan Cohen (Columbia University). One PhD and one postdoc will be hired by October 2024. (cont...)

Announcements (cont.)

The question of measuring the progression of aging has become crucial. Some think that it might be a predictor of health-related outcomes, others, that it would help determine if some drugs really are geroprotective, i.e., slow down the progression of aging. There is an intuitive idea that biological aging can be measured more accurately than with the progression of time, but this raises conceptual difficulties and the need for a concerted strategy. Is it possible to measure the progression of one and the same process across tissues? If not, how is it that we can frame such a project?

Fridolin Gross will lead the ANR project CELLTYPE. The project starts from the observation that recent developments, especially genomic single cell experiments, have made the concept of cell type, which has never been clearly defined, problematic.

The main aim is thus to clarify the concept. In particular, it posits two hypotheses: 1) Different, seemingly contradictory definitional criteria can be reconciled by explicating their conceptual relationships, suggesting a monistic notion of cell types as "functional kinds". 2) New findings do not lead to a conceptual revision, but to a revision of the factual knowledge about cell types. The project applies a combination of conceptual and historical analysis as well as an "empirical" approach in which scientists' underlying intuitions about cell types are systematically assessed through direct interaction and on the basis of surveys.

The project has a duration of 48 months and includes the employment of a PhD student, who is to be hired in fall 2024.

Publications

Pierrick Bourrat, William Godsoe, Pradeep Pillai, Tarik C. Gouhier, Werner Ulrich, Nicholas J. Gotelli, Matthijs van Veelen (2023). 'What is the price of using the Price equation in ecology?' *Oikos*: https://doi.org/10.1111/oik.10024.

Reports from recent PiBM events

PhilinBioMed Summer School – Bordeaux

Bordeaux's summer school in philosophy in biology and medicine took place this month, from June 12-16. For more details, visit: https://www.philinbiomed.org/event/summer-school-philinbiomed-2023/

As part of the 2023 PiBM Summer School, the PhilInBioMed Award was given to Elliott Sober.

The PhilInBioMed Award is given for an outstanding contribution to the advancement of biology or medicine through the use of philosophical and theoretical tools. To complement this recognition, the laureate received a prize of €5,000, which was awarded by the PhilInBioMed Institute in Bordeaux, France.

The PhilInBioMed Institute was pleased to announce Elliott Sober as the first laureate of the PhilInBioMed Award. Elliott is the Hans Reichenbach Professor and William F. Vilas Research Professor in the Department of Philosophy at University of Wisconsin-Madison, USA.

The title of Elliott Sober's Award Lecture, which he presented at the opening day of the PhilInBioMed Summer School, was: "Darwin on Group Selection and Phylogenetic Inference – Simpson's Paradox and the Law of Likelihood". For those who were unable to attend, or for any who wish to rewatch it, you can now find a video of the lecture here: https://www.philinbiomed.org/event/pibm-award-2023-sober/

Andrew J. Ewald, Ph.D. - Lecture on cancer biology at Bordeaux

Andrew is the Virginia DeAcetis Professor and Director, Department of Cell Biology and Director at the Giovanis Institute for Translational Cell Biology at Johns Hopkins Medical School. (cont...)

Reports from recent PiBM events (cont.)

The talk, given on June 20, was entitled: "Cellular strategies and molecular mechanisms driving breast cancer metastasis". Further details and the video of the talk can be found here: https://www.philinbiomed.org/event/andrew-ewald-cancer-invasion-and-metastasis/

Talking about PiBM

Thomas Pradeu interview in The Lonely Pipette

In June, Thomas Pradeu was a guest on The Lonely Pipette podcast. Thomas is a research director (CNRS) in philosophy of biology and codirects the Conceptual Biology and Medicine team within the ImmunoConcEpT laboratory at the University of Bordeaux, France. In episode 24 of the podcast, entitled "The untold power of philosophy", Thomas shares reflections on his journey navigating from philosophy to immunology and biology.

The Lonely Pipette is a podcast that invites guests to share tips about their life and careers in science. Each episode is a long-format conversation (often 45-60 minutes, in English) with researchers around the world about their career path, the way they run their lab and how they achieve life-work balance, among other topics. The hosts are Renaud Pourpre, a dynamic young science communicator and Jonathan Weitzman, a seasoned professor of genetics and epigenetics.

Jonathan has led many interdisciplinary projects bridging biology and philosophy and was an editor of *Le Dictionnaire Encyclopédique de l'identité* (Folio Essais, Gallimard 2020). The podcast began in 2020 and releases a monthly conversation with inspiring researchers who talk about their habits and recommendations. Most of the past guests have been biologists in academia, at various points in their careers, but Season 2 has branched out beyond biology and includes discussions with a chemist and researchers from industry. Thomas is the first philosopher to join The Lonely Pipette guest list. (cont...)

Talking about PiBM

In this episode, Thomas explains his views about the difference between 'philosophy IN biology' and 'philosophy OF biology'. He argues that philosophers should be brought in at the beginning of a project rather than at the end. He describes his model for embedding philosophers within a biology research environment. Thomas, Jonathan and Renaud have a lively discussion, sharing thoughts about personal successes and challenges. And Jonathan makes a desperate plea for a collaborator who is interested in the philosophy of parasitology.

Episodes of The Lonely Podcast, which carries the tagline "helping scientists do better science", have been downloaded more than 67 thousand times in 80 countries worldwide. Past episodes can be found on all popular podcast platforms and social media sites. Jonathan and Renaud welcome feedback from listeners and suggestions for future guests.

Here are various links to listen to the podcast on your favorite platform:
https://thelonelypipette.buzzsprout.com/share
To subscribe to the podcast newsletter : https://bit.ly/TLPNL
Twitter : https://twitter.com/LonelyPipette
Linkedin : https://www.linkedin.com/company/the-lonely-pipette

Interview with Ralph Adolphs by Thomas Pradeu

(Interviewed on December 8, 2022; transcribed on June 24, 2023)

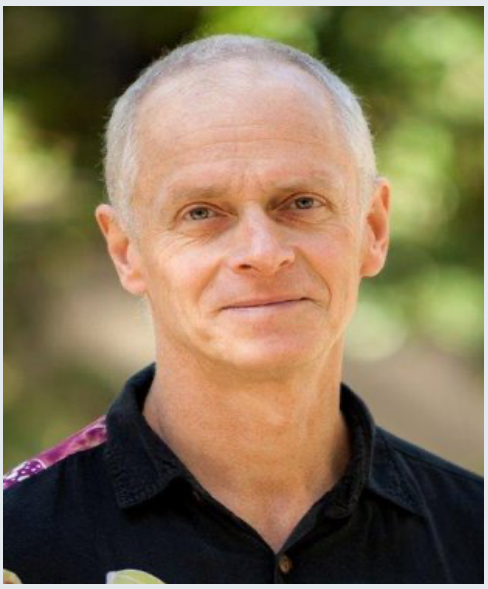

Thomas Pradeu: Today, it is my great pleasure to have a discussion with someone I admire a lot and who is also a friend of mine, Ralph Adolphs. Ralph Adolphs is Bren Professor of psychology, neuroscience and biology at Caltech in California There, Ralph is the Director of the Emotion and Social Cognition Lab, which investigates the neural underpinnings of human social behavior. One of the main questions that Ralph raises in his research is how people recognize, perceive, and process emotions.

Talking about PiBM (cont.)

With David Anderson, Ralph wrote a great book, entitled *The Neuroscience of Emotion: The New Synthesis* that was published in 2018. Ralph is also a member of our Philosophy in Biology and Medicine (PhilInBioMed) Network.

The aim of this interview is to explore with Ralph the role of conceptual and theoretical thinking as well as the role of philosophy in neuroscience.

Thank you, Ralph, for agreeing to do this interview. So, let's start with one of these very broad questions that we, philosophers, tend to ask. I think this is a good way to start the discussion. In current neuroscience, from about 1980s to today, who are the main scientists, and/or maybe the main books or the main papers that strike you as being very much involved in foundational thinking in your science, in the sense that they pay a lot of attention to conceptual aspects or theoretical aspects of the neuroscience?

Ralph Adolphs: That's a pretty big question. Of course, it goes back to well before the 1980s, and I think the main thing that happened in neuroscience is cognitive neuroscience and computational and theoretical neuroscience. Cognitive neuroscience is basically what I and my lab do.

It's hard to put dates on that. There were certainly people talking about foundational issues, psychologists and philosophers, primarily, well before then, but there wasn't that much known about the brain until modern methods came about. And then, since the 1950s, 1960s, 1970s, it's just kind of taken off. So, the field would basically be computational, theoretical and cognitive neuroscience. There would be a lot of people, so here at Caltech we have computation and neural systems as a graduate program which many other universities have as well, and some of the key players were people like, you know, Carver Mead, John Hopfield, Terry Sejnowski.

At least then, and to some extent also today, I think what distinguishes some of the main contributions is that these people didn't start in neuroscience.

(cont...)

Talking about PiBM (cont.)

They often started in physics or mathematics – you know, very quantitative sciences where you were used to writing down equations to describe the behavior of a system and trying to see how far you could go in applying that to the brain.

Again, it was limited a lot by data initially, and then we have more and more data now. There's functional imaging in humans. There are lots of dense recordings from cells in animals. And so given those data, the challenge is now can we put equations on those? Can we understand how the brain works in a way that ultimately, to paraphrase Feynman, we would be able to engineer one? Can we understand the human brain sufficiently that we could construct an artificial brain? That theme, which in the last five years is accelerating extremely rapidly, has also led to the marriage of artificial intelligence and machine learning with neuroscience.

So, it's an extremely rapidly evolving field that I think if you ask what's foundational and conceptual thinking in neuroscience, this is one of the main ones, though there are other conceptual strands as well. But I think the main approach is thinking of the brain as a computer. This is computational and theoretical neuroscience, with papers that go way back. One of the seminal papers (which certainly predates your date of 1980, but I think if you asked cognitive neuroscientists many would mention that paper) is the 1959 paper by Lettvin, McCulloch, Maturana, and Pitts, aptly titled "What the frog's eye tells the frog's brain". In this study, they were able to record from the frog's brain neurons and how they responded to visual stimuli like flies out there in the world. And ask the question: How does the frog 'know' that there is a fly so it can catch the fly? They found that there was already a lot of information at the level of the retina and the eye of the frog, there were cells that looked like feature detectors that were very specially designed to pick up, you know, little dots moving, like flies and so forth.

A bit later there was the Nobel Prize winning work of Torsten Wiesel and David Hubel, who recorded in cats to visual stimuli that they presented on the screen and asked, what do we see in the brain and the visual part of the brain when we do that?

(cont...)

Talking about PiBM (cont.)

And they characterized cells that had what they called 'receptive fields' that responded to certain visual stimuli in certain parts of the visual field. Then it really mushroomed out given these data and given that people are recording what happens in the brain when we show something out there in the world.

Computational and theoretical and cognitive neuroscience set it as their task to really try to figure out how does the brain represent the world, and the big part of that question, of course, is how does the brain generate conscious experience. This used to be a question that people relegated to philosophy until the 1980s, really probably when Pat Churchland published a book called *Neurophilosophy: Toward a Unified Science of the Mind-Brain* in 1986. And then lots of other neuroscientists and philosophers started investigating consciousness.

"How does the brain represent stimuli to guide action?" is one big set of questions. And then the harder question of how, why, in what way is that accompanied, at least in some cases, by conscious experience of the world?

There's a lot that one could say about that. Maybe to cut to one paper, I'd like to mention the article "The importance of mixed selectivity in complex cognitive tasks", by the senior author Stefano Fusi, published in *Nature* in 2013. It's a paper that really set a lot of current theoretical neuroscience. It shows the importance of mixed cell activity in complex cognitive tasks. As usual, this was offered by a bunch of people, but Stefano Fusi, a theoretical neuroscientist at Columbia University, is often credited with the theoretical conceptual thinking behind this paper. The gist of it is something like this: when people recorded from visual cortex from the optic tectum in the frog, they found neurons whose responses were easily interpretable; you had something on the screen, some dot, and this neuron responded just to that, and we could sort of tell a story for what is the function of this neuron. Its function is to detect flies and then tell the motor system of the frog to catch the fly. When people started recording other parts of the brain, in particular the prefrontal cortex (in the paper by Stefano Fusi, it was in the prefrontal cortex of monkeys), they found...

(cont...)

Talking about PiBM (cont.)

... neurons that responded to lots of things and their responses were not easily interpretable (much more similar in some ways to what you would find in the intermediate layers of a deep neural network nowadays). And the question was, how do we think about them? What do they do?

The upshot of the paper was that single neurons are not the way to think about how the brain works. For a number of reasons, but one reason is that the neurons further on that gets input and that ultimately are told by these neurons what kinds of representations or stimuli are out there are mappings from many, many thousands or tens of thousands of neurons onto these readout neurons. So any neuron listening in my brain to other neurons from which they get input gets input from thousands of other neurons, not just one. And so the question is, how does the neuron read out this population level, neural activity? The gist of this paper was that we need to record from many neurons in the brain and we need to understand a population level activity.

So, if the question is how does the brain represent a fly or an object or anything out there in the world, the answer is something along the lines that it constructs a very high dimensional representation of many thousands or tens of thousands of neurons, each of which responds in a mixed way in terms of their selectivity, to stimuli out there in the world. From that, the brain is able to read out many different things flexibly so that I can choose to recognize lots of different objects, classify them in different ways, in ways that are both very specific and very abstract.

It's a deep question, because right now when I'm looking at you on the Zoom, I'm able to represent specifics: you have the scene in the background, and you have the bridge and you have yourself and it's a very specific representation. On the other hand, I know that you're Thomas and you're a philosopher, etc. And that level of abstraction has to somehow be possible in the brain as well.

So this paper by Stefano Fusi is one along a line of papers that is really foundational in asking...

(cont...)

Talking about PiBM (cont.)

...Given we have the experimental data now and the methods we're able to record from all these neurons in the brain, how should we think about foundational questions about what does that tell the person and the rest of the brain, what does it represent? How could we read out information that this population of neurons represents about the world, and thinking about that conceptually and then coming up with tools to analyze those data, that would be one seminal thing.

Thomas Pradeu: **One thing I wanted to hear more about is to what extent you think that theoretical thinking in neuroscience has often been done by people coming from a different field, such as physics or mathematics, for instance.**

Ralph Adolphs: It certainly is the case. I mean, I think historically it's certainly been the case. So you know prior to lots of neuroscience methods and data being available, people like Donald Hebb or before that William James wrote about things that were about the brain, but there wasn't much known about the brain yet. So these were essentially psychologists. Psychologists were often studying the behavior of rats or descriptions of what people report in psychophysics experiments, but it was psychology, and those psychologists wanted to understand what it is that we could infer about the brain and what the brain represents given the behavior and the psychological data. One seminal paper there would be Edward Tolman's paper on cognitive maps in 1948 ("Cognitive maps in rats and men", *Psychological Review*). That paper started a whole line of research that's still extremely influential nowadays. Nobel prizes were given out for it, and there's lots of work in animals looking at cognitive maps in regions of the brain and lots of functional neuroimaging studies work. For instance, one could cite papers by people like Tim Behrens and colleagues that are very influential, with major big papers.

So, psychology I think had a big input and then philosophy. I mean, I think it's fair to say that, certainly for the neuroscience of consciousness, but also for cognitive neuroscience generally, the input from philosophers has been huge. It has basically been two strands, so one is probably of more relevance to the philosophers than the neuroscientists, and that's like Pat Churchland's book *Neurophilosophy*.
(cont...)

Talking about PiBM (cont.)

Neurophilosophy is basically telling philosophers to pay attention to neuroscience and incorporate that into things. But then philosophy of neuroscience is a big field in philosophy of science in general, because I think a lot of the questions about how do we think about causation, reduction and all of these questions that exist across philosophy of science, are especially acute in philosophy of neuroscience. And then there are specific people, like Peter Gärdenfors and others, who have thought hard about how do we think about a concept and how would that be instantiated? Then if you go to consciousness for sure all of the work in consciousness in neuroscience – and it's growing and growing, it's now a fairly respectable field, with societies, big grants, and journals dedicated to it – all of that really arose from, and is still mostly influenced by, philosophy, that is, philosophy of mind, and more specialized versions of it. Seminal books would be for instance David Chalmers' book *The Conscious Mind*. Chalmers is a philosopher, but I think in that book he was able to articulate very clearly, amongst many other people, what the problem is supposed to be. He made the distinction between the hard problem of consciousness and the easier problems, and that motivated lots of neuroscience work. So, historically, without question, philosophy and psychology have had a huge impact.

I think psychology may have a lesser impact today that it has had in the past. I think the input the infusion in terms of conceptual thinking in neuroscience for cognitive neuroscience and neuroscience of consciousness, it's really philosophy, philosophy of mind, and philosophy of science more broadly – I mean people like Carl Craver and others have spent a lot of time thinking about what counts as explanation, what is a mechanism, how do we think about the multiple levels in neuroscience that we need to account for.

Then, once computers came around, mathematics and physics had more of an impact. Decision making and decision neuroscience are big fields now that are very strongly predicated on mathematical models that come from machine learning, economics – fields where you could write down equations to understand how evidence is accumulated, decisions are made, in what sense they are rational with respect to some objectives.

(cont...)

Talking about PiBM (cont.)

So formalizing all of that has been a major infusion from several disciplines, but they're all formal disciplines: economic theory and mathematics have certainly provided pretty huge input there. It's been wedded nowadays very tightly with artificial intelligence, and the question of how do we design artificial systems that would make decisions, how does that process look like, what actually happens in the brain? Same thing for perception, obviously. How do we design objects like my cell phone that can recognize faces? Some of these artificial systems are extremely good at object recognition, including faces. Whenever I fly internationally now I have this global entry thing or whatever it is, I arrive in the United States immigration, I don't show my passport; I come in from a plane from Japan and they just look at my face and it's Ralph Adolphs. It's better than humans at recognizing individuals from their faces, so that's huge success. AI is very closely interacting with neuroscience. Many people, like James DiCarlo at MIT, are trying to make very strong parallels between what goes on in the deep neural network: Can we actually assign the kinds of representations we see in layers of a deep neural network that does face recognition to stages of processing in the human or the monkey brain? And the answer is yes, they actually bear a lot of similarity to one another. So the bottom line is that neuroscience is highly interdisciplinary and I think it's fair to say cognitive neuroscience is essentially interdisciplinary. It could not be done without input from other disciplines for sure.

Thomas Pradeu: **Would it be fair, or an oversimplification, to say that philosophers from this point of view have mainly brought conceptual clarification, while mathematicians and physicists have brought theories, predictions, or models?**

Ralph Adolphs: I think roughly that's true. I mean, by its nature, a lot of the philosophy has been more conceptual and clarifying the questions rather than providing algorithms for actually analyzing data or something like that. But again, take the work on, say, causation, which is often taken to be one of the key ingredients of explanation rather than prediction; that I would have be able to construct a causal model, and that would explain to me how the brain explains how a person sees.

(cont...)

Talking about PiBM (cont.)

I would need to have some causal model about how visual input causes things in the retina – boxes and arrows or wiring diagram. There's a lot of work on this in philosophy, but it merges into computer science.

In fact, one of the people that I collaborate with here at Caltech is in philosophy at Caltech, and has philosophical training but if you look at what he does, you would say it's computer science. He basically looks at causal discovery and algorithms, how from the statistics of data we could infer causal models, what kinds of assumptions are required, and so on and so forth.

Philosophy certainly contributes big conceptual questions like the nature of consciousness, and how to take that seriously and how to think about it deeply. And I think that's really important. And in part, it's because philosophers, at least that aspect of philosophy, can afford to be or takes the liberty of being detached from the data to some extent. And so philosophers have the freedom to really think deeply and pose very deep questions in a way that neuroscientists might find silly. They might say 'oh, I don't think about that, let's just go and look at the data and do the work'. But you need to step back from the data and ask really difficult questions and so that leaves you with questions, not answers, but I think that's important, as you need to have the questions before you can have the answers!

And then the next step, of course, for which philosophy is famous for, is more methodological. Philosophy doesn't just sort of step back and sit there pondering metaphysics, but it also has a method that helps to clarify the questions and provide some structure as to how you might find answers. I think the clarity for how to think about a problem is also something that philosophy has certainly contributed. For my own part, as you mentioned, we've interacted quite a bit. Moreover, I used to be on the McDonnell Project in philosophy and the neuroscience, run by Kathleen Akins back in the 1990s. We went to meetings on the West Coast of Vancouver Island each year and sat around and chatted about philosophical problems in neuroscience and these were all young people at the time, young philosophers thinking about the brain, Adina Roskies, Thomas Metzinger, lots of folks.

(cont...)

Talking about PiBM (cont.)

They are now well regarded senior philosophers working in neuroscience. And we did exactly the things that I just mentioned: we thought of the deep questions, ways to approach that. And it was just fun.

That ingredient, namely that philosophy should be fun, is important. If you're a graduate student or a postdoc in neuroscience, cognitive neuroscience for sure, you should certainly be reading and exposed to and interact with philosophers, because it gives you the fun of actually thinking about really hard questions with a lot of freedom and thinking about really foundational issues. How those foundational issues, then, translate into tangible progress in neuroscience is of course a very circuitous and difficult road.

There's another aspect which would also be considered foundational that I should mention that I guess is not from philosophy, but rather from within neuroscience. But it's very different and it's not fun. It's boring. I recently read a book by philosopher Michael Strevens, *The Knowledge Machine*. He paints a picture of science as very tedious; data and observation and actual results are what the scientific method is, and it always trumps everything else. And so science is hard, nature doesn't reveal herself to us very easily, we just need to crank through and see what the data tells us. And indeed, that's the thread that's been around certainly for the last ten years, and probably earlier than that. There are many people involved in that, but it's part of the so-called replication crisis that I suppose was most acute in psychology. People like Brian Nosek and many others started the Open Science framework. The basic idea applies to cognitive neuroscience in spades, and examples from cognitive neuroscience were often given. If you look at the literature on functional image brain imaging, fMRI, studies of human cognition, most of them are wrong. That's a fact. Most of them do not replicate. And when people realized this, it was hard to know what to make of it. How do you do science then, if most of the things published are wrong, have a huge bias, and so forth? People rose to the challenge and have risen to the challenge and said, well, the answer is...

(cont...)

Talking about PiBM (cont.)

... that you need to preregister studies, you need to have replications of studies and you need to really go through it very methodically and have very hard criteria that prevent you from conceptual thinking, prevent you from playing around and trying hypotheses or exploring. It's just something that you could imagine implementing in an algorithm that doesn't even need humans. It sounds extremely boring and in cognitive neuroscience one of the people that has had a big impact there is Russ Poldrack, a researcher now up at Stanford. So that's I think one domain that arose from within neuroscience, i.e., the realization of how difficult it is to actually have a cumulative science where at the end of the day you can really put confidence on your results. To do that you can't do it on your own, you need a whole community, and everybody has to be participating, in a very hard-nosed way.

That is a very different view than 'foundational thinking' about the problem of consciousness, where it it's much more creative.

I think marrying those two – sort of very hard-nosed approaches to how to get data and replication and get conclusions that will stand the test of time on the one hand and on the other hand thinking about hypotheses and trying to make sense of and understand how the brain works, which is a lot more fun – is one of the main challenges in cognitive neuroscience. But they're both foundational. One is foundational more in terms of the rigor of the methods and the other more foundational in terms of what it is that motivates us to do cognitive neuroscience. If all we had was the need to replicate and get data, it would be difficult to find people that want to do that. It's just not fun, right? So you need to have both. I think, which is what Michael Strevens stresses in his in his book that I just alluded to as well that neuroscientists are motivated by something that's in fact not, at the end of the day, how the science should be done.

But clearly you need to have the foundational conceptual thinking to come up with the questions in the first place. I mean otherwise the search space is infinite.

(cont...)

Talking about PiBM (cont.)

So you need to pose the questions and at the end of the day you need to have it to reflect on the data, once you know that they replicated, that they are sound, etc. to interpret what it is that they actually mean and come up with the next generation of questions. And I think this what to me makes cognitive neuroscience so exciting and so fun, because you really do have both of those ingredients.

Thomas Pradeu: **You mentioned that there is a role for philosophers in relation to neuroscience, but you also said something which I think is extremely interesting, which is that maybe we don't need the philosophers to do the philosophy, maybe this is something that could be integrated into the training of neuroscientists, and maybe in that sense philosophy would be better done by neuroscientists rather than by professional philosophers. So just, what is your view about this? And philosophy is done by neuroscientists themselves, what kind of philosophy are we talking about? What kind of philosophical training do you think would be needed to do this sort of more 'internal' philosophy by the neuroscientist themselves?**

Ralph Adolphs: As a field, if you step back, there is a lot of interaction between philosophy and neuroscience, but I would very strongly say that there should be much more interaction at the level of individual people and mechanisms to facilitate this, for example where a lab like my lab would host a philosophy graduate student. In the recent past, I had a student who did her thesis with Thomas Metzinger, she was here for six months, sort of a small visit and we had interactions, lab meetings. So, labs that do neuroscience hosting philosophers, especially young philosophers, or students doing their PhD in philosophy, that would be tremendously useful. All of this would start early, at the level of graduate students and postdocs. Once you had that, it would obviously translate up to when these people get old and are professors. But also the other way around: I think it should be a requirement that graduate students and postdocs in neuroscience should get courses in philosophy of science and philosophy of mind. It just seems like an obvious added ingredient that would not only make their work more fun, but would really allow them to hook up and link up with other people, and with the literature that's much broader than just their little particular niche.
(cont...)

Talking about PiBM (cont.)

So again, as a field, if you step back you would say there's a lot of influence between philosophy and cognitive neuroscience, and that's true, but it's not there yet at the level of actual individual mechanisms like labs, graduate programs actually making an effort so that people from philosophy or people from neuroscience could sort of swap over to the other side, and so I think that's extremely important.

Thomas Pradeu: **Don't you think that some of your colleagues might say that this is, at least in part, a loss of time?**

Ralph Adolphs: Some people would say that. I think again it would come down to the details of exactly which lab, which field and what are you going to be doing.

But I think the argument against that would be twofold. One is to just list the benefits and say that foundational thinking background in some of these conceptual issues and also the methodological rigor that philosophy, in fact, do bring a benefit to neuroscience. That's a pretty reasonable argument to make, certainly in cognitive neuroscience. If you're working on perception or representation or consciousness, you need to have philosophy in your background, or else you're just uneducated. So of course it depends to some extent on the field. The other argument, I think, would be something I alluded to earlier. It's broader than just saying it would be fun; it would also expose them to what I think needs to be defended, which is that a big part of neuroscience should just be wondering about big questions rather than being able to justify some application or curing some disease, et cetera. So the value of basic research really does pose these foundational questions. We're just curious about how does the brain work? How is it that I can perceive that people can think? How does language work? How do feelings work? All these are foundational basic questions and a value should be assigned to that. We have discussions often at my lab meetings about this issue.

That's the same thing that somebody would say with respect to art. You should be able to have an argument that there's an intrinsic value to just part of human nature is to be curious about ourselves and the world, and...

(cont...)

Talking about PiBM (cont.)

... think about deep, foundational issues without needing to say you know exactly what the application is, and I think that's really important and so philosophy would be exactly that. It's posing big, difficult questions that are part of what makes us human to puzzle about and to think about. We have that capacity, and I think it's part of our nature. Valuing that and taking time out to actually do that would also help. If you are in a neuroscience lab and all you're doing is focusing on methods and data, it can be difficult to raise these foundational questions, especially if all you hear is the purpose of this is to find a cure for Alzheimer's disease, or the purpose is to build something that's an extremely narrow applied goal. And of course it's well known that those applications arise from basic research as a side effect. But I think one wants to make a stronger argument that there's a value to just thinking about conceptual, foundational issues. That's part of who we are. It's the same answer to you know why have culture and why do art, independently of an application.

Philosophy values this and helps us teach how to think about those questions in a way that we're not guilty thinking about them. So I think that's really important.

Thomas Pradeu: **Two questions about your work on emotions. First, what has been for you, in that specific case, the value of clarifying the concept of emotion? Second, one thing that strikes me in your work is that comparative biology is important for you, so a lot of what you did was to say that emotions could be just in humans, but it's probably not the case and there's a way in which by being comparative, by examining what you call 'emotion primitives', we can gain a lot. This is something you explain very well in your book *The Neuroscience of Emotion: The New Synthesis*, with David Anderson. So, I would like to hear you about both conceptual clarification and the value of comparative biology, especially when we want to do the sort of foundational thinking/theoretical thinking that we mentioned before.**

Ralph Adolphs: That was a lot of fun to write that book. I wrote it primarily when I was in a four-month sabbatical in Kyoto, Japan, and...

(cont...)

Talking about PiBM (cont.)

... I worked together with the person whose office is immediately next door to mine here, David Anderson, who works on emotion but not in humans. He works in mice and in flies even, and so we came from very different species and different backgrounds. So, his background really for studying emotions was very mechanistic and very neurobiological, and my background was more infused with psychology.

If you broadly look in the literature on emotion, these two disciplines, or sub-disciplines, of emotion should be brought together, but are not. People working on emotion in animals and looking at neural circuits publish in completely different journals than people working on human emotions. Even if people working on human emotion do fMRI studies, most of their work is just psychology, and a lot of it is social psychology. It's published in psychology aligned journals.

So, bringing that together was tremendous fun. Both David Anderson and I firmly believe that's the only way, ultimately, to understand emotions. You do have to make comparisons across species. I think many people who think foundationally in neuroscience would subscribe to that theme. What we are after is something like fundamental principles in neuroscience, like how does the brain represent, or what are concepts? What is emotion? The answer to those questions should be broad enough that it's not just about modern adult humans or College students that I happened to study in the lab, but something that generalizes across species. Doing so gives us an answer to how it came about in the first place. So, yes, I think it's extremely important to have comparison across species and to some extent I haven't done this myself, but I do think it's very important. Also, comparisons across the lifespan are also important, so both evolution and development would be essential to give you a broad picture because things didn't come out of nowhere to understand questions such as: What were the building blocks throughout phylogeny? How did this do all up from birth on, to give you an answer on how something is actually created rather than just taking a single snapshot in an adult human brain when everything is all wired up and ready to go.

When you do this, you approach emotion conceptually by making contrasts and comparisons. (cont...)

Talking about PiBM (cont.)

You know there's not a necessary and sufficient condition type definition, and in fact nobody would expect that for emotion or any other concept. But you can certainly delimit it more and come up with operating characteristics and comparisons. So, the questions are obvious, and everybody would ask them. Are hunger or thirst or pain emotions? Can a dog feel emotions? Can an ant have emotions? And so you're already starting to try to delimit what emotions are, by contrast to these other things, and when you do that, you start coming up with criteria. So you might answer yes, because it has these particular properties or no, because it has some other properties.

I think that's basically the example that David Anderson and I went through in our book. We try to delimit what we think is a reasonable starting point as an iterative concept of emotion when we look across species. Then we ask ourselves how they are distinct from other kinds of internal states that we might not want to call emotions, and why, depending on operating characteristics. So we had a table in the book a lot of which is in agreement with what other people have written: emotions have a certain similarity structure of valence and intensity, they generalize, etc. That's just like a laundry list in a sense of what it is that they can do. But backing out from that then would essentially be a functional definition of emotions. In order to have all of these characteristics that we see that we think constitute emotions, what is it that you need to be able to do. That I think is the program in our book. Evolution across species came up with mechanisms for instantiating these internal states that we call emotions that have certain operating characteristics, they're ways of controlling behavior that are more flexible and generalizable than reflexes, but less so than just open thinking in humans. And the question is: What would you need in a system functionally in order to instantiate that?

So, that's the kind of program that we have there, but it does require abstraction and comparison across species. There's a colleague of mine here at Caltech, Markus Meister, who is also very synthetic and abstract and is thinking about how the brain works. He started a course maybe three years ago called "principles of neuroscience" where he tries to do exactly that.

(cont...)

Talking about PiBM (cont.)

What are the basic principles that distinguish neuroscience as a science from other aspects of biology or chemistry and physics? That all requires abstracting across species, trying to find generalizable principles of how nervous systems work, what is it that nervous systems do? Why did they evolve? What is it that we can find shared across species? And then of course, a part of that excercise is also saying what might be unique in the case of human nervous systems, but to get to that question you have to understand what it is that we have in common.

Thomas Pradeu: **So, overall, would you say that neuroscience is a field where foundational thinking is valued? Is it valued enough? Should it be valued more?**

Ralph Adolphs: It's valued, it depends on the sub discipline, but certainly in cognitive neuroscience and in theoretical computational neuroscience, it's definitely strongly valued.

And, yes, I think it should be valued even more, and in particular in terms of the details of how scientists go about their everyday business in particular in educating students. So like what we were saying before, neuroscience students should take philosophy courses, philosophy graduate students should be hosted in neuroscience labs, there should be more mechanisms for interaction between these disciplines even at the bottom level.

And then there is the threat I mentioned before. There's a strong emphasis now on data-driven atheoretical not hypothesis-driven work, lots of emphasis on replication and on big data. I could see that being a threat to at least some aspects of foundational thinking in that some people may say just let the data speak for themselves, everything should just be built on data, etc. To some extent, this is what you see if you look at some of the largest movements recently in neuroscience, like the Brain Initiative that we have here, with lots of funding. Although this has changed a little bit now, such grants have been really focused in the past on methods. How do we measure everything about the brain? How do we generate data? These are, of course, good aims, but the worry would be they might to some extent come at the expense of attention to foundational thinking. So we're poised now in neuroscience. (cont...)

Talking about PiBM (cont.)

We have a lot of methods now, we have dense neuronal recordings, we can record from thousands of neurons in a rat or a mouse, we have optogenetics for manipulating neuron activity very precisely in these animals. I think now we want to kind of come back and say, given these tools that we have now, we should actually take a breath and step back and return to foundational questions. And to some extent that's happening. There's a lot of work, including funding initiatives like from the brain initiative that are valuing more questions about how to analyze the data or how to think about the data.

So, it's dynamic, it's going in waves, but right now my feeling about neuroscience would be to indeed take a deep breath, take a look at all the methods that we have right now, we have more data and more methods than we know how to answer questions with them. We should actually put more of an emphasis on foundational thinking and return to big deep questions about how the brain can represent anything in the world, how representations can dynamically progress in time so as to constitute a train of thought or some kind of language of thought, and of course, ultimately, how does conscious experience arise from all of this? So, coming back and valuing foundational conceptual questions even more than we have in the past is something that will be really exciting to do over the next couple of years, because we have a lot more data and tools available than we did even five years ago.

Thomas Pradeu: **Ralph, thank you, that was a great discussion.**

Ralph Adolphs: Thank you, Thomas. You made me start my day here at Caltech on a happy positive note. I will be now in a foundational thinking mode for the rest of the day!

Looking ahead...

Phew! Hopefully that lengthy but delightful interview left you inspired and invigorated to advance and promote PiBM. Clearly, philosophy can and do contribute to science and (some) scientists are eager for the collaborations. Let's keep these discussions going as we pass through the summer/winter months!