Predictive Minds

Old Problems and New Challenges

Edited by
Manuel Curado
University of Minho, Portugal
Steven S. Gouveia
University of Porto, Portugal

Cognitive Science and Psychology

VERNON PRESS

www.vernonpress.com

In the Americas:
Vernon Press
1000 N West Street, Suite 1200
Wilmington, Delaware, 19801
United States

In the rest of the world:
Vernon Press
C/Sancti Espiritu 17,
Malaga, 29006
Spain

Cognitive Science and Psychology

Library of Congress Control Number: 2023941942

ISBN: 978-1-64889-743-6

Cover design by Vernon Press using elements designed by macrovector / Freepik and Freepik.

Table of Contents

List of Figures

Introduction

Manuel Curado
University of Minho, Portugal

Steven S. Gouveia[1]
University of Porto, Portugal

1. The Original Contribution of the Predictive Processing Model

The Predictive Processing Theory of Mind is a research project developed by philosophers, cognitive scientists and neuroscientists in order to understand an intermediary zone that lays between the activity of brains and the subjective experience. This approach has been described also as Bayesian brain, predictive coding, active inference, just to name a few current designations (cf. Friston, 2005; Clark, 2013; Hohwy, 2013). The novelty of this new way of thinking about action, perception and mental representation needs to be explained. It is not an investigation of brain cells and neural networks, nor is it an investigation of the contents of the minds of agents. These two realms characterize different types of agents, whether biological or made from other materials, and are undoubtedly important. However, these realms do not seem to have contact zones with each other. Indeed, what could a biological object with weight and internal structure like the brain have to do with the subjective experience of human beings? In each of these realms it would be impossible to find signs of the existence of the other realm. When we walk on a beach early in the morning and see footprints in the sand, we may conclude that someone has been to the beach before us. Unfortunately, no brain reveals the mental life of its owner, and it is also true that no human being has subjective experiences of his own brain. You may feel pain or pleasure, but you will never feel your neurons, and vice versa. In analogy with other physical systems, similar problems are found. Looking around, people see water and think that the world has a reality called liquidity. The truth is

[1] This work was carried out initially by the grant Project UIDB/00683/2020 funded by FCT at the Centre for Philosophical and Humanistic Studies of the Universidade Católica Portuguesa, Braga, Portugal, and was finalized under the CEEC Individual Fellowship 2022.02527.CEECIND by FCT at the Mind, Language, and Action Group of the Institute of Philosophy, University of Porto, Portugal.

that an object called liquidity does not exist anywhere. If someone wants to buy only liquidity, and not objects that are liquid, he will not find a store that can sell it. There is no such store. Liquidity seems to be a high-level property that only exists for intermediate-sized beings, neither too small nor too big. If human beings were the size of atoms, when looking at a glass of water they would only see hydrogen and oxygen in a certain proportion and, obviously, they would not see a third body called liquidity anywhere. The reality of liquidity and the reality of water molecules cannot be called into question, but the relationship between the two needs to be explained, and this explanation is more complex than it seems, because the statements that affirm the existence of liquids are unintelligible without observers who have nothing to do with either hydrogen or oxygen. The innovative project of Predictive Processing, in turn, seeks to exhaustively explore the intermediate zone between the two realms of the brain and mental life, an area with enough complexity to justify the need to also consider it as autonomous, a third realm with its own organization. In the case of water liquidity, the connection between the two sides to the subject requires a whole theorizing process to explain it, a result that took many centuries to achieve. By the same logic, in order to get an explanation of the inner details of the connection between sides so radically different, we need creative concepts and time.

What we know already about this third realm is that the brain is a prediction machine with the function of creating hypotheses about the causes of sensory signals and of making predictions of possible future sensory signals. Moreover, the internal models of the world that are created this way are constantly challenged by incorporating the errors of the previous models into new models. From this point of view, the work of the brain could be described as a process of making predictions about the upcoming sensory data based on its best current models of the causes of those data. This idea is still not sufficiently explained and needs a lot of applied research and extensive theorizing work. The problems are many. It was stated above that the two realms of mind and brain do not reveal signs of the existence of each other, just as water molecules do not reveal the existence of liquidity, nor does liquidity reveal the existence of hydrogen and oxygen. The problems do not go away with the assumption that there is a third instance that connects the two sides of the problem, or, resorting to the analogy again, the two realms. Incidentally, the assumption of a third instance further complicates the problems. In the area of contact between the brain and this third instance, no signs of either side are identified. In turn, in the contact zone between this third instance and the complexity of the mental life of agents such as human beings, the problem reiterates itself. In the case of water, explaining the existence of the property of liquidity requires time and the search for the best possible explanation, an abductive mode of reasoning that is never quite finished. In the case of Predictive Processing, assuming the existence of a sub-

personal level that cannot be confused with either the mind or the brain is, at this moment, the best explanation we have about the relationship between agents and the situations in which they find themselves.

It has often been repeated that the Kantian idea of reversing the traditional approach to the mind is similar to what Copernicus did regarding the place of the Earth and the Sun in the Solar System, a Copernican revolution that "assume that objects must conform to our cognition" (Kant, 1998/1787: 100, sec. B xvi). This change in perspective constitutes the core of the theoretical proposal of the Predictive Processing paradigm: the perceptual experience of the world is decisively influenced by cognitive systems that generate models of reality. The reception of stimuli coming from the outside is not a merely passive process, but a very active one. Neither the stimulation that agents experience at a given moment nor unforeseen situations immediately affect the model of reality they take as a reference. The details of this process have not yet been fully inventoried, and it is therefore difficult to have a clear idea of the scope of this new way of understanding this intermediate zone between the materiality of brain cells and the splendor of the phenomenology of the mental life of agents. It is possible that a much sought after and much needed theory of everything about mental life is being constructed. This theory would have the mission of explaining all aspects of the mental life and experience of human beings, certainly, but also of other types of agents.

Let us focus on a practical example of what is at stake with the Predictive Processing framework and its explanation of every single mental feature of the mind.

The first mental feature — and perhaps the one that was in the original point of depart of the reflection about the scope and the goals of the Predictive Processing framework — is the full set of our perceptions of the external world. Most theories of perception have to give a plausible answer to the problem of explaining how it is possible for the brain to process sensory data from the world the result of which is considered a mental representation caused by autonomous objects of the real world. As Hohwy puts it, "[t]he problem of perception is the problem of using the effects — that is, the sensory data that is all the brain has access to — to figure out the causes" (Hohwy, 2013: 13). Even philosophers of mind that do not work within the Predictive Processing paradigm, advancing other conjectures about the ultimate nature of the conscious mind, have emphasized the importance of studying the nature of the connection with external objects. Chalmers, for example, shows that central processing only finds itself in an information space, because it can no longer access external objects and the very physical processes that underlie perception (Chalmers, 1996: 290).

Predictive Processing is a theory that will try to answer to the problem of perception with a somehow surprising answer, far from the common-sense

view: perception is a kind of controlled hallucination created by the way our brain functions. Contrary to the traditional views about perception, that consider the mental process as a passive and stimulus-driven process, the Predictive Processing framework challenges that view by arguing that perception should be considered as having a constructive nature, being active and influencing constantly sensory data received from the external world.

Imagine that one particular human brain is inside a skull in a similar way that a person is stuck inside a room with no direct access to the external environment. There is, however, a significant difference in the two situations. The person in the room has nothing but signals to conjecture about what is outside. She can be deceived, as any recording can lead her to think about a certain subject; she may also misread the signals. In any case, what the person inside the room has is just a pale image of what possibly exists outside. If she hears the sounds of falling rain, she won't get wet; if it's a tropical sun, she won't tan. This person knows that what is inside the room is not the real thing because signals are not things. In contrast, the brain inside the skull, despite not having direct contact with external objects, nor, indeed, with sensory organs that have been stimulated by external factors, creates, from the signals that reach it, a reinforced reality. In this case, the rain indeed wets and the sun indeed tans. As the person in the room only has conjectural images of what is outside, she cannot have subjective experiences of the alleged referents of those images. The notion of reality that she has is very poor. Differently, the reinforced reality created by the brain inside the skull is surprisingly stable, given the rapid changes that occur in perception, whether due to environmental factors or due to the action of the human body.

The results of the brain working inside the skull are so extraordinary that they cannot be bettered: they have the unmistakable stamp of what is real. When the person leaves the room where she is, her experience of reality is significantly increased. The same does not happen with the experience of reality produced by the brain. There is no rational way to increase the reality of the experience of things. The lights may be more intense, the sounds may be louder and the blows received may be stronger, but none of this will alter the degree of reality of the experience. Dim light and bright light will be equally real, and there is no known alternative way to create anything remotely resembling the unmistakable touch of reality. Like an anchor that holds a ship to the bottom of the sea, the unmistakable mark of reality constitutes a stable support that the changes that occur on the surface of the sea cannot break. The analogy is imperfect, because, in the case of the ship, sufficiently large alterations can cause the anchor to separate from the ship, but, in the case of sensorial alterations, the notion of reality cannot be broken. Perception organs may be damaged or even destroyed, but this does not alter the notion of reality. The transfiguration of a mere internal representation

made by the brain in the darkness of the skull into a reality insusceptible of improvement is precisely the type of stability that deserves to be studied.

In literary or mythical terms, there is a large number of stories in which the protagonists make representations so perfect that they mistakenly take them for realities. The person inside the room may have a high-definition television with her, but she will not confuse the perfection of the images with the strange experience called reality. St. Anselm of Canterbury's ontological argument can also be invoked to help understand that the best of illusions cannot be confused with the worst of realities. There is an insurmountable difference between thinking of a being with all perfections without the perfection of existence and thinking of a being with all perfections including the perfection of existence. This difference is not only of interest for the exceptional theoretical case of proving the existence of God. The simplest act of perception is equally exceptional because, more than a mere representation of something, it offers the conviction that this something exists. We could have a set of thousand photographs of an object, without ever confusing the first with the second. A thousand images are not an object and all the perfections without the perfection of existence are worth little. Now, it is precisely the point at which the metamorphosis of mere representation into a sensation of reality takes place that deserves to be understood. Flickering Galateas transform themselves into a beautiful maiden with whom Pygmalion wants to marry. To understand this is the main objective of Predictive Processing. It appears to be a merely cognitive matter (whatever that is), but it is, in fact, a way of illuminating the zero degree of reality construction.

How does the brain perform so well, if that's the case? Rephrasing the question: if the brain doesn't have direct access to the world, how can it be so successful in our daily interactions with a dynamic external world? The answer of the Predictive Processing framework is to claim that the "locked" brain has to do probabilistic inferences about the causes of the sensory data by using available sensory evidence and previous experiences that will serve as a basis for creating a predictive model of the world.

Here is an intuitive example of what the Predictive Processing means by claiming that the brain is a prediction machine. Please focus on the following image:

Figure 0.1. Dots and spots.

If the reader focuses on the image (see Figure 0.1.), assuming that you never encountered that image before, you will probably think that it doesn't make any sense and that there is nothing in those dots, spots, shapes and shadows that is accurate or represent a known entity from the world. Is it a photo of skin marks? Is it a pattern of the Martian landscape? No one can tell. Is impossible to name or pinpoint a referent of the image.

Your brain is absolutely clueless about what this illustration should represent and that's because the predictive model of your perceptual experience in this moment doesn't match anything that you have previously experienced and, beyond that, the sensory inputs don't make sense either.

However, this is where the Predictive Processing Theory seems to be plausible. Check the same image below, now with a simple additional line. The new detail will make you have a sense of what the image does represent:

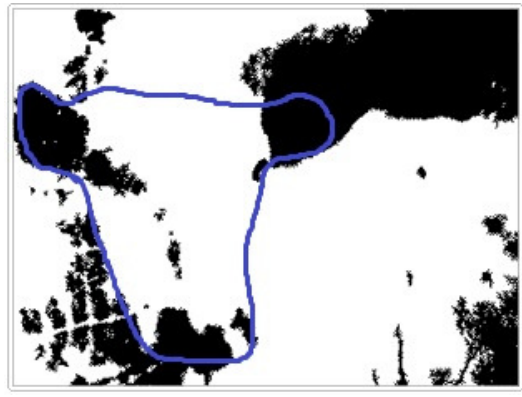

Figure 0.2. PP example.

With the help of the added line, your brain can now make an accurate perception of the image above. The second image (Figure 0.2) allows your brain to update the previous model and to create a new model where your brain can now detect the sketchy shape of the head of a cow. If you go back to the previous image (Figure 0.1.) — recall that that didn't make sense the first time you looked at it —, you could verify that that same image will right now look like a cow. Even without the added new line an average percipient is unable to look at that image and *not* seeing a cow since the internal model was updated with the new version processed in Figure 0.2. The new reality was stabilized and reinforced and it becomes more difficult to change the second image than to change the first one with the contribution of the line that was added to it.

It is an interesting question to try to know if the line that was added could be replaced, for instance, by a verbal description or a gesture. The first time we look at the image, a friend might ask us to imagine a line shaped like a cow's head. That friend could even run a bare finger across the image, tracing an outline. A photocopy of the image could be cut out with scissors to form the shape of a cow's head. These three alternative means (the words of suggestion, the finger gesture, and the paper cutout) seem also to play a top-down role in helping the brain shape meaningless visual impressions. A skilled magician manages to get the attention of the spectators who watch his show by organizing in a certain way the vision that these people have of what they are seeing on stage. Similarly, the three alternative means contribute to the emergence of a particular form. As soon as this happens, the internal model is updated and it becomes almost impossible to see again the initial image without the new organization.

This is just one brief example of the sort of knowledge that the Predictive Processing framework can offer regarding the nature and functions of the embodied mind. The core concept of this framework is inspirational for many future lines of inquiry. For instance, how far could it be used in order to explain aspects of the mind more difficult than perception, if indeed there is something more difficult than the creation of the experience of stable reality? As we unconsciously use spatial metaphors to designate the parts of the mind, the relationship of the mind to the brain, and even mental states, it seems to us that perception is "at the bottom" and that other functions, such as reason, are "at the top." It is possible that nothing in mental life is as difficult to explain as perception, because only perception allows the experience of reality. A memory or a belief pales in comparison to perception. However, this line of inquiry has many other associated problems. If a sub-personal system is able to create the experience of reality, the full understanding of what is at stake in this process may allow that same capacity to be applied, for example,

to memories. In this case, the construction of the sensation of reality would not be limited to the present time. A second possible line of thought questions the relation between sub-personal processes and formal aspects of nature. For instance, is there any structured similarity or common constraint between the role played by the anticipatory internal mental model and non-mental processes in nature (e.g. entropy, evolutionary processes, imago in the metamorphosis of insects and morphogenesis)? Inspired by parallel projects of mapping neuronal connections and by literary (e.g. Marcel Proust, Virginia Woolf) and philosophical descriptions of the current of consciousness (e.g. William James, Edmund Husserl), Predictive Processing is also challenged to exhaustively map possible geometries of the dynamics of these sub-personal processes. The question to be asked in this case would be the following: Can we build a geometric model of the main structure of the third realm that seems to exist between the realm of biology (the brain) and the phenomenological realm (the subjective mind)? Assuming that this hypothetical third realm could be fully mapped, would this imply that the second realm of human mental experience could also be mapped? And, of course, if a major breakthrough in neurosciences provides in the remote future a full cartography of the work done by neurons and glia, as well as of their connections, would this result promote the full mapping of the other two realms? Many exploratory questions like these could be formulated due to the richness of Predictive Processing.

Wealth, did we say? Or maybe it's better to say poverty, because the line that was added to Figure 0.1. still hides many secrets. The amount of information in this Figure is greater than in Figure 0.2. This means that the ability to get a stable image of the cow's head has paid the high price of losing information. Many consequences derive from this impoverishment, and it is not easy to explain the meaning of what happened in the passage from the first to the second figure. Every moment of human life reveals the same constraints. All the words we use end up fulfilling the function of the line that was added to the collection of points and shapes: they all imply a paradoxical impoverishment of available information, paradoxical because it is this impoverishment that allows human action. Generalizing the results of the examples, the totality of the mind seems to suck into a bottomless abyss the excesses of information from the environment. There is no action by agents without each of them carrying out the bad deal of becoming poorer in one of the parameters of reality.

2. New Developments

The main goal of this book will be to critically analyze and discuss different aspects related to this new framework. The first chapter, by Wanja Wiese and

Thomas Metzinger, offers a general overview of the core concepts focusing on the hierarchical predictive coding, the way precision is mediated, the prediction error minimization, and finally how the proposal accounts for predictive control. This introductory chapter aims to be an introduction that will enable readers to follow the subsequent chapters of the collection making them familiar with some of the main general principles of Predictive Processing and some of its key concepts.

Next, Arthur Schwaninger will examine the metaphysical implications of the Predictive Processing framework by concluding that it supports a view by which reality is the result of the subject's history of experiences and actions. Following this metaphysical investigation, Manuel Curado will present an essay on the metaphysical constraints of Predictive Processing by reflecting on the intellectual ambition that Predictive Processing manifests, contextualizing the hype in the long history of failures of scientific and philosophical explanations of the functioning of brain and mind. Believing that nothing less than an ultimate theory of mind is at stake, he takes into consideration the structural similarity between the creation of information by the brain, the structured levels of cognition and the wide aspects of the physical organization of the world. Curado confronts this ambitious agenda with difficult metaphysical questions, namely the evaluation of difference and sameness in human affairs and the high price that is paid by the alleged advantage of sub-personal processes of prediction.

Moving on from Metaphysics to Emotion, the next three chapters will present three reflections on how Predictive Processing and specific emotions can be related in different settings. First, Dina Mendonça shows how the new framework requires the integration of the complex multilayered richness of the emotional landscape in order to be a unified theory for understanding the mind, including the affective character of experience. The chapter concludes by promoting the idea that the understanding of the richness of human emotions and its role in cognition implies an education of the emotions themselves. Next, José Araya will claim that emotions can be understood not as a form of perceptual inference — as claimed by some Predictive Processing theorists — but as forms of regulation based on active inference. To close this topic, Mark Miller and Jelena Markovic will investigate how affective neuroscience can cast some light on the emotion-cognition interactions in the predictive brain.

The next three chapters will show how diverse and useful the Predictive Processing framework can be as a theory by focusing on mental diseases. First, Jorge Gonçalves explores this problematic by showing that, if the theory of the Predictive Processing aspires to give an overall explanation of the human mind, the motivational factors of delusion should be considered.

Next, João Fonseca will focus on the relationships between the self and body illusions by claiming that that core-self ought to be understood as a hierarchically layered predictive processing structure. Finally, Valtteri Arstila, Jarno Tuominen and Susanne Uusitalo explain how gambling disorder can be seen from the point of view of Predictive Processing, showing how gambling behaviour could be considered rational. This new interpretation allows a more informed assessment and personalized treatment decisions that could potentiate the defeat of the gambling disorder.

To close the book, three more chapters will be presented. Matthew Sims will show how Predictive Processing can be a limited theory regarding direct perception in contrast to an ecological psychology. Next, Jona Förster will argue that Predictive Processing describes a teleofunctional computational mechanism. In the last chapter, Miguel Pais-Vieira, Manuel Curado and Nicolás Lori will present a method to study how the qualitative aspects of information could be related to quantitative aspects and vice-versa. Several authors allude to the internal informational space, but no one has determined the exact number of states that can be represented in the human brain (or as far as the argument goes, by any other type of brain) and how those states can be related with time and other physical constraints. The mathematical method that is described should be seen as a tool to help present and future researchers pursuing the high-priced goal of a complete theory of information in the brain. This will be relevant for the Predictive Processing assumption that the human brain represents information and, in order to do so, needs to use levels of cellular energy, neural networks and other physical constraints. The ghost-like nature of information and the wide-ranging scope of definitions that that concept inspired created the illusion that something intangible is happening inside our skulls. However, without cellular processes and networks it would be impossible to represent information. The mathematics of this third realm is still in its first stages. The final chapter is, therefore, a contribution to that goal.

As can be seen, each chapter in this book proposes new lines of reflection and possible contexts for applying the paradigm. There is conceptual exploration, there is intellectual subversion through critical thinking taken to the limit of what can be thought, and there are proposals for mathematical methods to face what is possibly the greatest challenge of the Predictive Processing paradigm: the nature of information in the brain.

After reading this book, we hope that the reader will have a broader notion of what kind of novel and interesting research can be made from a Predictive Processing point of view. The next years will be absolutely relevant to determine whether Predictive Processing can really become a Theory of Everything regarding the mental or if it can only provide specific explanations

of particular mental phenomena. This book is certainly a step into the critical consideration that needs to be made in order to achieve a sound answer to this doubt.

References

Chalmers, D. J. (1996) *The Conscious Mind: In Search of a Fundamental Theory.* New York: Oxford University Press.

Clark, A. (2013) 'Whatever next? Predictive Brains, Situated Agents, and the Future of Cognitive Science.' *Behavioral and Brain Sciences,* 36 (3): 181–204. doi. org/10.1017/S0140525X12000477.

Friston, K. (2005) 'A Theory of Cortical Responses.' *Philosophical Transactions of the Royal Society B: Biological Sciences,* 360 (1456): 815-836. doi.org/ 10.1098/ rstb.2005.1622.

Hohwy, J. (2013) *The Predictive Mind.* New York: Oxford University Press.

Kant, I. (1998) *Critique of Pure Reason.* Translated by Paul Guyer and Allen W. Wood. New York: Cambridge University Press.

PAGES MISSING

FROM THIS FREE SAMPLE

Contributors

Arthur Schwaninger is a junior researcher at the Department of Philosophy at the University of Zurich, Switzerland. He has both a Master's degree from the University of Cambridge and from ETH Zurich. Before working in philosophy, he conducted research at the Human Brain Project, modeling and performing computer simulations of biological neural networks in the visual cortex. He published numerous articles in the domain of computational science as well as in philosophy and is member of the Digital Society Initiative in Zurich. In recent years, his research has focused on alternative naturalistic approaches to metaphysics from a cognitive science perspective.

Dina Mendonça (Ph.D. University of South Carolina, United States, 2003) is a research member of IFILNOVA and works on Philosophy of Emotions and Philosophy for Children at the New University of Lisbon, FCSH NOVA (Universidade Nova de Lisboa).

Her current research is focused on the application the Situated Approach to Emotions, which takes emotions as dynamic and active situational occurrences (Mendonça 2012) to Ethics and Decision-Making. This research of Deweyan inspiration aims at elaborating a critical interpretation of the philosophy of emotions clarifying different methodological and philosophical approaches to emotions, as well as identifying the key issues and problems for emotion and further complexities of our emotional world (Mendonça 2013, 2014, 2017). She is the author of several papers on emotion theory and, in addition, promotes and creates original material for application of philosophy to all schooling stages, and as an aid in the creative processes.

Jarno Tuominen is a teacher and a senior researcher at the Department of Psychology and Speech-Language Pathology at the University of Turku, Finland. His PhD in Psychology concentrated on consciousness research, and more specifically on the social simulation theory of dreaming. His current research interests additionally include behavior change, altered states of consciousness and mentalization.

Jelena Markovic a PhD candidate in Philosophy at the University of British Columbia, specializing in philosophy of mind and cognitive science. Her doctoral dissertation is on transformative experience, specifically transformative

experiences that we do not choose to undergo such as grief, accident, and illness. In addition to articles on transformative experience, she has also published articles on the cognitive science and philosophy of attention, including on affect-biased attention, hypnosis, and meditative states.

João Fonseca is a researcher at the Nova Institute of Philosophy (IFILNOVA, New University of Lisbon). His main areas of research are Philosophy of Science, Philosophy of Neuroscience and Cognitive Science. He is the author of several papers and is co-editor (with Jorge Gonçalves) of Philosophical Perspectives on the Self (2015). Currently, he is developing a semantic/model theoretic framework for behavioral neuroscience in order to clarify topics such as the generalization of neuroscientific explanations and the relations between neuroscience and psychology. He is also interested in the impact some topics from Philosophy of Mind (the Hard-Problem or Mental Causation) have in the neuroscientific paradigm on Fear Conditioning.

Jona Förster is a PhD student of cognitive neuroscience at Humboldt-Universität zu Berlin, Faculty of Philosophy, Berlin School of Mind and Brain, and at Freie Universität Berlin, Department of Education and Psychology, Neurocomputation and Neuroimaging Unit. He is interested in philosophy of science as well as in cognitive neuroscience. His doctoral research is on the neural correlates of somatosensory awareness. He is currently delving deeper into the active inference framework, and hopes to receive his swimming badge soon.

Jorge Gonçalves has been a postdoctoral researcher at IFILNOVA – FCSH (New University of Lisbon) for twelve years. He has a Doctorate, a Masters, and a Degree in Philosophy and a further Degree in Psychology. His PhD thesis *The Place of Consciousness in Nature*. His current research area is the Philosophy of Psychiatry, with a special focus on the "concept of mental illness," "mind-body problem and psychiatry," "theories of delusion," and "the problem of the unconscious." Jorge has worked as a Clinical Psychologist for twelve years. He is the author of several papers and editor of three books. He was PI of the funded project "Cognitive Foundations of the Self." He was also part of a funded project on Philosophy of Movies.

José M. Araya is Associate Researcher and Director of the Department of Systemics at IFICC (Chile). He is also Psychology Lecturer at Universidad de Talca (Chile). He obtained his PhD at the University of Edinburgh. His work

focuses on the philosophical and social implications of current research in cognitive and affective sciences.

Manuel Curado is Corresponding Member to the Brazilian Academy of Philosophy, Professor of Philosophy at the University of Minho, namely in the areas of Philosophy of Mind and Philosophy in Portugal, and obtained several other titles, namely that of National Defense Auditor (Ministry of Defense, Lisbon) and Senior Management Course for Public Administration (University of Minho/INA). He is doctor cum laude (University of Salamanca, Spain) having previously obtained a master's degree in Philosophy (New University of Lisbon). He has delivered over 300 conferences in many countries, from Moscow, Russia (MGLU and MGIMO), to Brasilia, Brazil (UnB), and was Erasmus Professor in Padua, Italy. He is the author of several books, including the first critical edition of an eighteenth-century manuscript written by a fellow of the Royal Society of London, the Portuguese Jewish physician Dr. Isaac Sequeira Samuda, MD (*The Viriadas of Dr. Samuda*, IUC/Coimbra University Press, 2014). Along with his philosophical activities, he has a literary taste for promoting the edition of Portuguese-language classics, such as D. Francisco Manuel de Melo (*First Treaty of Cabala*, Círculo de Leitores, 2018), Teixeira de Queiroz (*Complete Works*, Imprensa Nacional-Casa da Moeda, 2020-2023) and João da Rocha (*The Wisdom of Patience*, C.M.V.C., 2021, and the critical edition of *Anguishes*), because he glimpses in them signs of perennial wisdom. He re-edited Rear-Admiral José Maria Dantas Pereira's *Memoir of a Pasigraphy Project* (1800), having promoted since the 90s the knowledge of this pioneer of the utopia of the perfect language and universal machine translation. He re-edited the works on philosophical psychology and logic by Edmundo Curvelo, and published his manuscripts (*Complete Works of Edmundo Curvelo*, Fundação Calouste Gulbenkian, 2013; *A Portuguese Genius: Edmundo Curvelo (1913-1954)*, Universidade de Coimbra, 2013). He was member to several ethics committees, and expressed his bioethical thinking in the book *Health and Cyborgs: Caring in the Biotechnological Era* (Edições Esgotadas, 2019). He recently co-edited *Perception, Cognition, and Aesthetics* (Routledge, 2019) and *The Philosophy & Science of Predictive Processing* (Bloomsbury, 2020).

Mark Miller is a Research Fellow at the Monash Centre for Consciousness and Contemplative Studies, Monash University (Australia). He obtained his PhD at the University of Edinburgh. His research explores what recent advances in computational neuroscience can tell us about happiness and well-being, and what it means to live well in our increasingly technologically-mediated world.

Matthew Sims obtained his PhD in philosophy at the University of Edinburgh in 2020 under the supervision of Dave Ward, Andy Clark, and Julian Kiverstein. Sims is currently an Alexander von Humboldt postdoctoral fellow at Ruhr University Bochum where he researches the evolution of goal directed behavior, collective intelligence, and basal cognition.

Miguel Pais-Vieira is an Assistant Professor at the University of Aveiro (Portugal) working in the field of neuroscience. His studies focus on brain-machine interfaces for the treatment of neurological diseases, but also on basic aspects of neurophysiology. His main goal is to understand the nervous system in its full complexity. He has, since 2019, been collaborating with Physicist Nicolás Lori and the Philosopher Manuel Curado in a joint multidisciplinary attempt to understand information processing in the brain.

Nicolás Lori is a Researcher at the University of Minho (Portugal) working in the fields of physics, neuroscience and computer science. His studies focus on Magnetic Resonance Imaging data processing, information in quantum mechanics and the ontologies of computational processes. His main goal is to understand the role of information in its full complexity. He has, since 2019, been collaborating with Neuroscientist Miguel Pais-Vieira and the Philosopher Manuel Curado in a joint multidisciplinary attempt to understand information processing in the brain.

Steven S. Gouveia holds a PhD in Philosophy from the University of Minho (Portugal) and is currently leading a 6-year research project on the Ethics of Artificial Intelligence in Medicine at the Mind, Language and Action Group of the Institute of Philosophy of the University of Porto (Portugal), being also an Honorary Professor of the Faculty of Medicine Andrés Bello (Chile). He has published and edited 13 academic books on interdisciplinary topics. He also organized several online courses with the participation of prominent scholars such as Noam Chomsky, David Chalmers, Sir Roger Penrose or Anil Seth. He hosted and produced the documentary "The Age of Artificial Intelligence". More information: stevensgouveia.weebly.com.

Susanne Uusitalo is a Collegium Postdoctoral Researcher at the Turku Institute for Advanced Studies and the Department of Philosophy, University of Turku, Finland. Her research interests are in applied philosophy, applied ethics and research ethics. Currently, her research focus is on issues concerning the ethics of addiction research and treatment, especially in

gambling disorder, and conceptual commitments of agency in addiction, substance use and gambling.

Thomas Metzinger (*1958 in Frankfurt am Main, Germany) is currently Professor of Theoretical Philosophy at the Johannes Gutenberg-Universität Mainz and an Adjunct Fellow at the Frankfurt Institute for Advanced Study (FIAS). Metzinger is past president of the German Cognitive Science Society (2005-2007) and of the Association for the Scientific Study of Consciousness (2009-2011). His focus of research lies in analytical philosophy of mind and cognitive science, as well as in connections between ethics, philosophy of mind and anthropology. In the English language, he has edited two collections on consciousness (*Conscious Experience*, Paderborn: Mentis & Thorverton, UK: Imprint Academic, 1995; *Neural Correlates of Consciousness*, Cambridge, MA: MIT Press, 2000) and published one major scientific monograph developing a comprehensive, interdisciplinary theory about consciousness, the phenomenal self, and the first-person perspective (*Being No One: The Self-Model Theory of Subjectivity*, Cambridge, MA: MIT Press, 2003). In 2009, he published a popular book, which addresses a wider audience and discusses the ethical, cultural and social consequences of consciousness research (*The Ego Tunnel – The Science of the Mind and the Myth of the Self*, New York: Basic Books). An important recent Open Access collection (2015) is Open MIND at http://www.open-mind.net, see also http://predictive-mind.net (2017) for a follow-up.

Valtteri Arstila is a Collegium Researcher at the Turku Institute for Advanced Studies and the Department of Philosophy, University of Turku, Finland. His research interests are mainly in the philosophy of mind and philosophy of psychology. Currently, his research focus is on issues concerning the philosophy and psychology of subjective time, the philosophical theories of color, cognitive penetration of perceptual experiences, and the prerequisites of strong artificial intelligence.

Wanja Wiese studied philosophy and mathematics at Johannes Gutenberg University Mainz, where he also obtained his PhD in philosophy. A revised version of his dissertation was published by MIT Press (*Experienced Wholeness*, 2018). In 2017, Wanja Wiese co-edited the collection *Philosophy and Predictive Processing* (together with Thomas Metzinger), which is freely available at https://predictive-mind.net. Wiese is currently a postdoctoral lecture and assistant researcher at Ruhr University Bochum. His research centers on consciousness, mental representation, philosophy of cognitive

science, and philosophy of artificial intelligence. He is editor-in-chief of "Philosophy and the Mind Sciences" (https://philosophymindscience.org), together with Sascha Fink (Magdeburg) and Jennifer Windt (Monash). *Philosophy and the Mind Sciences* is a peer-reviewed, not-for-profit open-access journal that is free for authors and readers.

Index

D

E

F

G

H

I

J

K

L

M

N

O

P

Q

R

S

T

U

V

W