2018 Rijeka

UNIVERSITY OF RIJEKA
FACULTY OF HUMANITIES
AND SOCIAL SCIENCES
10-13 SEPTEMBER 2018

Book of abstracts

(Updated 05.09.2018)

Contents

Key note lectures

José Luis Bermúdez (Texas A&M University)

Framing, rationality, and self-control

Abstract: Is it irrational to be influenced by how outcomes are framed? Almost everybody thinks so. Yet, experimental and theoretical discussions of framing have ignored the most interesting situations in which framing occurs. These are when an agent or decision-maker consciously and simultaneously considers a single outcome or scenario under two different frames and evaluates it differently in each frame. From the perspective of standard theories of rational decision-making and rational choice, such clashes of frames are not supposed to happen – and if they do occur they are paradigms of irrationality. The standard view, I claim, gets things almost completely wrong. There is nothing intrinsically irrational about seeing the world in conflicting frames. In fact, actively engaging competing frames can be a powerful force for rationally engaging with problems that classical theories of rationality are unable to tackle. This paper defends this basic idea and shows how it sheds light on psychological and decision-theoretic discussions of self-control.

Robyn Carston (University College London)

Lexical Creativity, Sense Conventions and the Lexicon

Abstract: Speakers can use language creatively to express a new ad hoc sense or to coin a new word, and hearers are usually able to grasp the new sense or word by virtue of their inferential pragmatic capacities. Such new senses or words may be transient innovations or may become conventional and stable components of the lexicon. Taking an interdisciplinary perspective (drawing on work in philosophy, psychology and linguistics), I will argue that (a) such pragmatically-based sense conventions are the essence of the phenomenon of (cross-categorial) polysemy, (b) conventionalised senses are stored in a communicational lexicon which is distinct from the lexicon of the narrow language faculty, and (c) contrary to some recent views in the philosophy of language, sense conventions, although hugely enabling, are not essential to the explanation of successful communication.

Michael Graziano (Princeton University)

Consciousness Engineered

Abstract: What is consciousness? How can a brain, a mere collection of neurons, create it, and can we artificially create it? In my lab we are developing a theoretical and experimental approach to these questions that we call the Attention Schema Theory (AST). The goal of the AST is to understand consciousness well enough to engineer it – to one day be able to build artificial, conscious machines. AST seeks to explain how an information-processing machine could act the way people do, insisting it has consciousness, describing consciousness in the ways that we do, and attributing similar properties to

others. AST is a theory of how a machine insists it is more than a machine, even though it is not. The theory begins with our ability to attribute awareness to others. The human brain has a complex circuitry that allows it to be socially intelligent. One function of this circuitry is to attribute a state of awareness to others: to build the construct that person Y is aware of thing X. The brain, in effect, builds a simplified model or simulation of how another machine is processing information. The semi-magical, mystical properties that people often attribute to awareness may be a result of the inevitable inaccuracies in those simplified models. In the theory, the machinery that attributes awareness to others also attributes the same property to oneself. This type of modeling serves a specific use: it helps to predict the behavior of other people, and to predict and control one's own behavior. I will discuss the basic framework of AST and some of the evidence from human behavior and brain imaging that supports this approach.

L.A. Paul (Yale University)

Transforming the self

Abstract: Going to war, having a baby, being spiritually reborn, betraying your lover, emigrating to a new country—all of these are experiences that can transform you. By transforming you, they replace your old self with a new self, and in the process, they restructure the nature and meaning of your life. I'll discuss the nature of transformative experience, and show how exploring the epistemic structure of the self-change involved can help us to understand the special and distinctive ways that new experiences can form and change us, how we use them to construct our lives, and how this relates to how we make life choices, both big and small.

Invited symposia

Agency, Normativity and the Social Self (Organizer: Kristina Musholt, University of Leipzig)

Heidi Maibom. (University of Cincinnati); Symposium on Agency, Normativity, and the Social Self.

Abstract: Western philosophy of mind has been dominated by a certain view of self-knowledge according to which there is a problem concerning the existence of other minds. We know, by means of infallible introspection, that we have minds, but we can never quite be sure that others do. This line of thinking is rather difficult to square with the fact that we are social animals and with what we know of child development and human psychology. My talk concerns the ways in which we are socially constituted. The first thing to note is that our understanding of ourselves as having conscious minds is the result of meeting other minds. This is a psychological truth, even though a conceptual argument can be made to that effect as well. Understanding our own minds as minds only arises in the intersubjective context. It presupposes the existence of other minds. What is central to understanding our minds as minds is perspective. Being a conscious being means having a certain perspective on the world. This may seem unexciting, relatively speaking, and a somewhat abstract and metaphysical point. But it carries over to very practical aspects of life. Specifically, understanding what we are doing requires us to be able to take other perspectives on our actions than our own. It requires us to be able to see what we do as others might, in particular those affected by our actions. Knowing what we do is not something we can know purely by means of introspection.

Daniel Yudkin (Yale University), Annayah Prosser (Yale University) and Molly Crockett (Yale University).

Moral Transformation at Multi-Day Mass Gatherings

Abstract: Many people experience an event at some point in their lives that results in a profound change to their moral values. However, the processes underlying such "transformative experiences" are not fully understood. In this talk, I present an investigation into the psychological of transformative experiences involving participants of so-called "transformational" mass gatherings in the US and UK. In a series of studies including 1,305 on-site and 10,568 online participants, my collaborators and I found evidence of significant changes across a range of behavioral and attitudinal measures, including social inclusivity, moral evaluation, connectedness to others, and generosity. Transformation at "gift economy" festivals (i.e. those barring money) in particular corresponded with increased sensitivity to the "warm glow" of prosocial behavior in moral judgment. Follow-up studies in the weeks, months, and years after attendance demonstrated the persistence of some (but not all) of these changes over time. Overall, this research provides a richer empirical picture of the psychology of prosocial transformation, with a particular emphasis on the unique psychological properties of money-free environments and the value of "warm glow" giving.

Mattia Gallotti (London School of Economics)

Co-Cognition and Social Discourse

Human life develops and flourishes in a world of shared habits and perspectives. Research into social cognition now suggests that new avenues to knowledge and experience of things become available to individuals when they see things 'together' and act upon shared representations. In my talk, I'll discuss this theoretical development in the context of a recent contribution to the philosophical debate by Jane Heal (2013). Heal argued that considerations about the relevance of acts of 'co-cognition' provides a wider context in which to analyse, not just the social dimension of self, but also fundamental aspects of mentality and agency. I will focus in particular on the transformative nature of co-cognition – i.e. the 'affordances' of social interactions – and its role in enabling individuals to do more, or different things, by having their perspectives expanded as they align minds and bodies in interaction.

Normative and descriptive approaches to human reasoning (Organizer: Shira Elqayam, De Montfort University)

In this companion symposium to José Luis Bermúdez's keynote address, we bring together recent philosophical, psychological and pragmatic contributions to the study of the interplay between the 'is' (descriptive) and the 'ought' (normative) in human rationality.

Cognitive Success – Gerhard Schurz will present a model of cognitive success that combines ecological validity and applicability.

Conditionals and Multiple Norm Conflicts – Niels Skovgaard Olsen addresses the relationship between normative and descriptive aspects by measuring norms as psychological profiles.

Pragmatic Persuasion - Michaela Wänke will present a pragmatics-based approach to inference in persuasive communication.

The symposium will be concluded with a roundtable discussion led by the keynote speaker, José Luis Bermúdez.

Niels Skovgaard Olsen (University of Göttingen). Conditionals and Multiple Norm Conflicts

Abstract: Whenever multiple, competing norms (e.g. classical logic vs. non-monotonic logic, probability theory vs. ranking theory) can be applied to a given reasoning task, the problem of arbitrating among these norms arises.

In Elqayam and Evans (2011), this problem was, inter alia, used to argue that the Psychology of Reasoning should start being a purely descriptive discipline and eschew its normative theorizing. However, this is a view that has been met with heavy resistance both in the commentaries to Elqayam and Evans (2011) and the essays collected in Elqayam and Over (2016).

The purpose of this talk is to develop an experimental paradigm for addressing the situation, where we have multiple norms that deliver conflicting verdicts on whether the performance of the participants counts as correct or incorrect. Employing Steinberger's (draft) conceptual distinction between directive,

appraising, and evaluative uses of a formal system, we hereby attempt to address the abovementioned problem of arbitration as it arises in an appraising setting.

To illustrate, the conflict between the Suppositional Theory of Conditionals and a local, inferentialist theory of conditionals is used since these theories deliver opposing instructions for how to assign probabilities to conditionals (Skovgaard-Olsen, Singmann, and Klauer, 2016a, 2016b). Recently, and-to-if inferences have received much attention in the psychology of reasoning, because they were identified in Douven (2015) as diagnostic for experimentally distinguishing between the Suppositional Theory of Conditionals and Inferentialism. Through the present experiments, we seek to establish profiles of the participants to examine individual variation in their interpretation of indicative conditionals. Next we confront the participants with the uncertain and-to-if inference task to examine the extent to which participants that are classified according to different interpretations of conditionals are capable of reasoning correctly according to the given profile they have been assigned.

References

Douven, I. (2015). The Epistemology of Indicative Conditionals: Formal and Empirical Approaches. Cambridge: Cambridge University Press.

Elqayam, S. and Evans, J. St. B. T. (2011). Subtracting "ought" from "is": descriptivism versus normativism in the study of human thinking. Behavioral and Brain Sciences, 34, 233-90.

Elqayam, S. and Over, D. (eds.) (2016). From is to ought: The place of normative models in the study of human thought. Lausanne: Frontiers Media.

Skovgaard-Olsen, N., Singmann, H., and Klauer, K. C. (2016a). The Relevance Effect and Conditionals. Cognition, 150, 26-36.

Skovgaard-Olsen, N., Singmann, H., and Klauer, K. C. (2016b). Relevance and Reason Relations, Cognitive Science, 1-14.

Steinberger, F. (draft). Three ways in which logic might be normative. URL = http://floriansteinberger.weebly.com/research.html

Gerhard Schurz (Düsseldorf University). Cognitive Success

Abstract: I argue that the rationality of systems of reasoning should not be evaluated in terms of normative standards that are based on a priori intuitions but, rather, on their cognitive success. Cognitive success can be defined and gauged in terms of two factors: ecological validity (the system's validity in conditions in which it is applicable) and the system's applicability (the scope of conditions under which it can be applied). I show that prominent systems of reasoning — deductive reasoning, Bayesian reasoning, and uncertain conditionals — perform rather differently on these two factors. Finally, I demonstrate that conceptualizing rationality according to its cognitive success offers a new perspective on the time-honored relationship between the descriptive and the normative in psychology and philosophy.

Michaela Wänke (Manheim University). Pragmatic Persuasion

Abstract: Although social psychology defines persuasion as communication it has nevertheless neglected communicative aspects when studying the dynamics of persuasion. Based on research on social

communication (McCann & Higgins, 1992) and conversational logic (Grice, 1975; for a review see Wänke, 2007), the present research aims to fill this gap and to complement our understanding of persuasion by a pragmatic perspective that takes into account a recipient's inference of why information is communicated. According to conversational logic recipients may expect that communicated information comes with a "guarantee of relevance" (Sperber & Wilson, 1986), namely relevance for the purpose of the communication (Grice, 1975). If the purpose is persuasion, recipients may expect the presented information to hold potentially persuasive implications. In a nutshell, the pragmatic persuasion perspective holds that because receivers expect persuasive communication to convey information that is relevant to the goal of persuading them, they are likely to infer arguments in favor of the persuasion goal from the presented information. In turn, receivers may become persuaded by these inferences. In other words, simply being presented as part of a persuasive message can make information persuasive. I will present several studies that illustrate the implications of this perspective in the context of mass communication.

Polysemy (Organizer: Ingrid Lossius Falkum; Center for the Study of Mind in Nature, Oslo)

Ekaterini Klepousniotou (University of Leeds). Psycholinguistic evidence for the processing and representation of polysemy

Abstract: Having as a basis the linguistic definitions and categorization of lexical ambiguity into homonymy and polysemy, experimental evidence will be presented to support this distinction and the psychological reality of linguistic polysemy. Data from a series of studies using a variety of experimental methodologies will provide converging evidence for differential processing and representation in the mental lexicon not only between homonymy and polysemy, but potentially also for the different subtypes within polysemy. Implications for models of lexical ambiguity processing will be discussed.

Louise McNally (Universitat Pompeu Fabra). Polysemy: A window on grammar

Abstract: In this talk, based on ongoing work with Alexandra Spalek, I show how a comparison of patterns of verbal polysemy in English vs. Spanish sheds light on the organization of the verbal lexicon and the distribution across the verb root and its morphosyntax of the components of meaning that categorize eventuality participants vs. those that categorize other aspects of eventualities. Extensions of meaning from the concrete to the abstract domain show striking patterns of both similarity and contrast across the two languages. These patterns suggest that grammar shapes our categorization of eventualities in ways that we may not be aware of and highlight the deeply intertwined relationship between "grammatical" and "root" components of the linguistic sign.

Francois Recanati (Institut Jean Nicod). What polysemy teaches: meaning vs senses

Abstract: Whereas two homonymous expressions (e.g. 'bank' and 'bank') are different expressions, with the same phonological realization but distinct meanings, a polyseme is a single expression, i.e. a semantic as well as a phonological unit. It has one meaning, which should not be confused with the separate senses which it contributes in context. Different ways of thinking of that unitary meaning (Ruhl's and Langacker's) will be discussed.

The evolution of consciousness (Organizer: Liz Irvine, Cardiff University)

Jonathan Birch (London School of Economics). Can a Single Neuron Support Subjective Experience?

Abstract: Although there will always be disagreement about when exactly the "lights came on", we can make progress towards understanding the evolution of consciousness by understanding the evolution of sensorimotor integration. Recent work on the nematode worm C. elegans is illuminating in this respect. C. elegans is able to steer smoothly towards the source of an attractive odour, guided by an internal representation of the direction of the odour gradient in relation to its body. Remarkably, the integration of sensory and motor signals supporting this behaviour occurs within a single neuron. This gives a tantalising glimpse of what the earliest forms of subjectivity might have been like, and offers some support for deep gradualism about the origins of subjectivity. It also poses a challenge to vision-centred accounts of the evolution of consciousness, which tie the origin of subjective perceptual experience to the appearance of the first eyes.

Janicke Nordgreen (Norwegian University of Life Sciences). What research on pain and emotions in non-mammalian species can bring to the debate over animal consciousness

Abstract: After more than a decade of research on fish pain and nociception, the subject is still source of a great deal of controversy, with the question of whether fish are conscious being at the center of the debate. I will present different lines of argument both in favour of and against consciousness in non-human animals, and I will discuss whether research on animal pain and emotions can further our understanding of consciousness, with examples from three non-mammalian species; rainbow trout (Oncorhynchus mykiss), common goldfish (Carassius auratus) and laying hens (Gallus gallus domesticus). While experiments that target nociceptive responses directly might not be the best tools to help us understand an animal's conscious experience, the debate about the interpretation of such experiments and the attempts to define animal pain are valuable as starting points for research on consciousness in non-human animals.

Clint Perry (University of London). How foresight might support the behavioural flexibility of arthropods

Abstract: The small size of arthropods' brains has traditionally been thought to be a limitation to sophisticated forms of cognition and capable of only specific sets of pre-programmed behavioral routines. But evidence of cognitive abilities once thought to be exclusive to larger brains are more frequently being revealed in a variety of insects and other invertebrates. Locusts, grasshoppers, flies and dragonflies have been shown to use internal models of the surrounding world to tailor their actions adaptively to predict the imminent future. Honeybees and orb-weaving spiders appear to possess a "master plan" of their home constructions. Jumping spiders seem to pre-plan their routes indicating they understand the problems they face and actions necessary to obtain prey. Bumblebees and ants utilize non-natural objects as types of tools to solve problems in a manner that suggests an awareness of the desired outcome. I will describe some of these works and suggest that it might be simpler for these small brains to operate with cognitive tools such as foresight and awareness rather than to have specific behavioral routines for each environmental situation they come across. I will describe other works suggesting that bees possess basic forms of emotion and metacognition, discuss how these findings speak to the possibility of consciousness in these small brains, and suggest that what we learn will ultimately help us determine how consciousness has evolved and provide us with vital information to understand consciousness in all brains.

Liz Irvine. Discourse eliminativism and gradualism about consciousness

Abstract: In this talk I'll argue that discourse eliminativism about the concept of consciousness (defended in Irvine 2012) may be compatible with, and lend additional support to, a multi-dimensional gradualist view about (animal) consciousness of the sort sketched by Godfrey-Smith (2017). In particular, I'll argue that both of these views can be developed such that they challenge 'all-or-nothing' intuitions about

consciousness: that an organism is either capable of being conscious or it is not. Instead, I claim that giving a description of the multi-dimensional behavioural and cognitive 'package' associated with an organism is all that can be done, such that there is no further question to be asked of whether the organism is 'really' conscious. I claim that while counter-intuitive, there are biological and methodological reasons to take such a stance seriously.

Accepted papers and posters

Aina, Laura (Pompeu Fabra University)

Lexical modulation via distributional semantics: a general framework for modelling word meaning in context

Abstract: The content that a word introduces in a context often varies from situation to situation: this property of the lexicon encompasses a series of different phenomena and it is broadly referred to as polysemy (Falkum and Vicente, 2015). Such an ambiguity is resolved in communication via the process of lexical modulation, which allows deriving the speaker-intended meaning assigned to a word, given the word itself and its context of use (Cruse, 1986). With variation from theory to theory, the inputs to such process is taken to consist of, on one side, some context-independent information associated with the specific word (e.g., a concept, a collection of past uses etc.) and, on the other, contextual information associated with the situation, which may include both linguistic (e.g., previous discourse) and extra-linguistic content (e.g., perceptual information about the environment).

We propose a general framework for developing formal models of lexical modulation, employing the computational methods of Distributional Semantics (Lenci, 2008; Baroni, 2016). We argue for the appropriateness of adopting these to obtain rich and fine-grained representations of the inputs to lexical modulation, namely a word and its context. Within Distributional Semantics, words are represented by means of abstractions over their contexts of use given substantial amounts of linguistic and perceptual data, in the form of vectors of continuous values. With related methodologies, representations of contexts of the same format can be derived both for linguistic chunks of varying size (e.g. from phrases to entire discourses) and for situated information from perceptual modalities (e.g., an image representing the environment of communication). We regard this type of inputs to constitute a flexible starting point to design models of lexical modulations that accommodate hypotheses of interest. On one side, distributional representations of words can be taken as approximations of context-independent information associated with a word, that we argue to be in principle compatible with different views on word meaning. On the other, the methods employed in Distributional Semantics to represent contexts allow incorporating different types of situated information. Crucially, the functions that can be defined over these inputs are continuous in nature: they are for this reason apt to capture the nuanced and dynamic nature of lexical modulation. In the paper, we provide theoretical arguments and empirical results from the literature supporting this proposal. We aim at creating new synergies between theoretical and computational research on Polysemy by showing that Distributional Semantics has the potential to be a precious toolbox for linguists and philosophers of language: it can indeed facilitate the development and validation of theories by grounding them in large-scale data.

References:
Cruse, A. 1986. *Lexical semantics*. Cambridge University Press.
Lenci, A. 2008. Distributional semantics in linguistic and cognitive research. *Italian journal of linguistics*.

Baroni, M. 2016. Grounding distributional semantics in the visual world. *Language and Linguistics Compass*.
Lossius Falkum, I. and Vicente, A. 2015. Polysemy: Current perspectives and approaches. *Lingua.*

Aliffi, Matilde (University of Birmingham)

Emotions and Epistemic Responsibility

Abstract: In the first act of Shakespeare's tragedy 'King Lear', the King, looking for flattery, asks his daughters to express their love for him. When the third-born Cordelia turns down his request for flattery, the King reacts in a wave of anger and pride, and he vehemently disclaims his parental care for Cordelia. King Lear's emotion of anger seems utterly disproportionate and even irrational. There is nothing extremely offensive in Cordelia's speech that justifies such tremendous rage. More precisely, the King's anger appears to be epistemically irrational: it is presenting information that is not well supported by the evidence. Let us assume that King Lear's emotion is epistemically irrational. Is the King epistemically responsible for his anger?
Like King Lear, we often have emotions that seem irrational. In this poster, I am going to explore to which extent subjects are epistemically responsible for the emotions they experience. The idea that we are epistemically responsible for the emotions we feel seems prima facie appealing. We blame people that are misled by their emotions in their reasoning. If emotions are skewing our epistemic landscape (Goldie, 2000), we also seem to be partly responsible for the distorting effect of the emotions on our reasoning. Conversely, we also praise a competent ability to have correct emotions and to use profitably the information that emotions give to us for reasoning purposes.
I will argue that epistemic responsibility is grounded on the subject's ability to control their emotional states. Drawing on the extant discussion on the normativity of beliefs, such as (Olson, 2015; Alston, 1988) I distinguish three types of control of emotional states: synchronic, diachronic and fine-tuning control. I then develop an account of the degree of responsibility that the subjects enjoy for their emotions.

References
Alston, W. 1988. The Deontological Conception of Epistemic Justification. Philosophical Perspectives, 2, 257.
Goldie, P. 2000. The emotions: a philosophical exploration, Oxford, Clarendon Press.
Olson, D. 2015. A Case for Epistemic Agency. Logos & Episteme, 6, 449-474.

Alsmith, Adrian (Jean Nicod Institute)

Sensations in remote touching

Abstract: In remote touching, one perceives through touch the properties of an object that is not in contact with one's body. It is sometimes claimed that in remote touching, one experiences remote sensations, viz. sensations at locations of non-bodily contact: contact between two perceived objects, neither of which is a body-part. This paper aims to motivate denial of this claim by considering three views of remote touching. The first – the distal view – serves well to explain remote touch phenomena without appeal to remote sensations. The second and third – the illusion view and the extension view – are both apt to serve as a basis for claiming that there can be remote sensations. I will outline the distal view by describing how it explains psychophysical research often taken to demonstrate the existence of remote sensations. The

illusion view and the extension view might independently motivate positing such sensations; however, as I hope to show, each of these face significant issues.

Anderson, Curt (Heinrich-Heine-Universität Düsseldorf)

Of old couples and important committees: accessing members of groups using modifiers

Abstract: Adjectival modification of group nouns such as committee and couple show that group nouns have a complex lexical semantics; groups differ in the accessibility of members within the group, as well as in whether an adjective targets a property of the group or its members. We analyze group nouns using frames (Petersen, 2007) to model their lexical semantics, and expand the semantic ontology to include social entities. We show how types of group nouns differ in their conceptual structure and modification potential.

Asma, Lieke (Munich School for Philosophy)

Unconscious processes and conscious control

Abstract: Psychological experiments seem to show that conscious processes play a much less prominent role in what we do than we might have expected, even when it comes to higher mental processes (for an overview see, e.g., Bargh & Chartrand, 1999; Bargh & Ferguson, 2000; Caruso, 2012; Custers & Aarts, 2010; Dijksterhuis & Aarts, 2010). These studies have often been interpreted as showing that explanations of what we do in terms of external triggers and unconscious processes can replace a substantial amount of explanations in terms of conscious states and processes, for example intentions. The suggestion is that because of this, we lack conscious control over what we do. However, it is often not made explicit what conscious control exactly is, and it is unclear whether and to what extent agents actually lack it.
In this paper I aim to fill this gap and develop a notion of conscious control in connection to intentional action (see Shepherd, 2015). I argue that it is not at all clear that agents lack conscious control because of unconscious influences. Furthermore, I point out that in every situation in which we intentionally act there are also a lot of accurate descriptions of what we do under which we lack this kind of conscious control, but (1) often this lack of control is irrelevant because we do not care about the event that is not controlled, and (2) this is for an important part independent of whether what we do was (to a large extent) caused by unconscious processes. I suggest that the relationship between conscious control and unconscious causes is not as strong as is often assumed.

Bárány, Tibor (Budapest University of Technology and Economics)

"Simple Sentences": Conversational Implicatures or Ironical Utterances?

Abstract: In the literature it is commonly supposed that the 'symptoms' of linguistic intensionality—failure of substitutivity of coreferential terms or predicates, failure of quantifier exportation, failure of existential generalization, and availability of nonspecific readings—can be explained by the occurrence of certain expressions in the sentence. These expressions (propositional attitude verbs, modal expressions, certain adverbial modifiers, etc.) create intensional context: it is their syntactic/semantic scope within which the linguistic elements evoke our anti-substitution intuitions and intuitions about the lack of existential impact or about non-specificity. However, so-called 'simple sentences' (analyzed first by Jennifer M. Saul in 1997) do not contain intensional expressions, yet co-referring names cannot be substitutable in them salva veritate. (For example: "Clark Kent went into the phone booth, and Superman came out"—"Clark Kent

went into the phone booth, and Clark Kent came out".)

Semantic accounts of the phenomenon want to show that the names in question are, contrary to what we might think initially, in fact not coextensive (Pitt 2001; Moore 1999; Moore 2000); or the logical forms of simple sentences contain hidden indexical elements which block the substitution of co-referring names (Forbes 1997; 1999; 2006). Pragmatic accounts hold that our anti-substitution intuitions should be thought of as products of some purely pragmatic mechanisms—for example sentences which differ only in respect of co-referring names trigger different conversational implicatures (Barber 2000). According to 'the mistaken evaluation account' provided by David Braun and Jennifer M. Saul (2002) and Saul (2007) our intuitive judgments about the truth values of simple sentences that we get by substitution inference are incorrect, but what we can hold responsible for the mistaken evaluations of truth values are some complex psychological or cognitive processes.

In my presentation I offer a pragmatic explanation to the phenomenon of 'simple sentences' which is based on the so-called pretense account of verbal irony (cf. Currie 2006; 2011). Uttering a 'simple sentence' the speaker makes as if she would not know that the names in question are co-referring, while she takes for granted that the members of the audience surely know this—and that they are fully aware of the fact that the speaker knows this as well. The person who speaks ironically adopts simultaneously two (or more) epistemically (or otherwise) incompatible perspectives; uttering 'simple sentences' is only one way, yet very intricate way to speak ironically. A major advantage of my account is that by applying it to the phenomenon of 'simple sentences' we can withdraw our everyday anti-substitution intuitions (unlike Saul's psychological explanation or Barber's pragmatic account) without making dubious assumptions about the conventional use of proper names (Pitt, Moore) or the logical form (Forbes) of simple sentences.

Barlassina, Luca (University of Sheffield)

Embodied cognitive reuse: Use in moderation—but with feeling

Abstract: Alvin Goldman has recently developed a new approach to cognition: moderate embodied cognitive science (Goldman 2012, 2014). The approach is moderate in that it retains one of the tenets of classical cognitive science, namely, that cognition is a computational process over mental representations. It is embodied because it rejects another pillar of classicism: modularity--the idea that the mind is to a large extent composed of functionally specialised cognitive systems, each of which deploys proprietary mental representations. Instead, Goldman proposes that a significant chunk of cognition is based upon embodied cognitive reuse--the exploitation of embodied representations/cognitive systems for non-primary functions. For example, the motor system and motor representations, whose primary function is to control movement, are hypothesised to be reused for action recognition, language understanding, visual perception, etc.

I have good and bad news for Goldman. Bad news: he does not provide any convincing argument for the claim that embodied cognitive reuse is a pervasive aspect of cognition. Good news: when plugged into a by-and-large classical cognitive architecture, embodied cognitive reuse does a wonderful job in explaining at least one important cognitive phenomenon: emotion. Hence the title.

I proceed as follows. First, I criticise Goldman's master argument for the generality of embodied cognitive reuse. The argument starts from Anderson's (2010) neural reuse theory, according to which neural circuits established for one cognitive function are routinely reused to perform different cognitive functions. On this

basis, Goldman concludes that a large number of cognitive functions are underpinned by the reuse of embodied cognitive systems/representations, rather than by dedicated modules/representations. I show that this argument is a non-sequitur. It rests on a confusion between explanatory levels: neural and cognitive. Even if the neural reuse theory were true, nothing would follow about cognitive reuse.

Second, I build a new theory of emotion upon embodied cognitive reuse. My theory, I maintain, illuminates the cognitive underpinnings of emotions, accounts for a large swath of empirical evidence, and suggests new research questions. A toy example should give you the gist of it. You are afraid of a snake. This is how your emotion came about. The sight of the snake caused fear-typical bodily changes in you. The latter activated #H#, an embodied representation reused for the secondary function of representing the property being harmful. The tokening of #H# initiated a process of causal reasoning, which aimed to individuate the cause of your bodily changes. The process outputted #that snake#, and this representation was filled in #__ CAUSE H#. The resulting mental representation #that snake CAUSE H# constitutes your fear. It represents that the snake is a source of harm.

Here is the take-home message. Classicists can rest easy: modularity is here to stay. But Goldman should be satisfied as well: embodied cognitive reuse is an important tool. Classicists should add it to their theoretical toolbox. If they don't, they risk pushing emotions outside their reach.

Barnden, John (University of Birmingham)

Consciousness, Causal Oomph, and Meta-Oomph

Abstract: In this talk I argue, on the one hand, in support of a radical claim about the relationship between phenomenal consciousness and causation, and, on the other hand, for the need to focus more attention on an aspect of causation itself. The claims are in the context of an anti-Humean and dynamicist stance, in which there is an irreducible dynamism or causal "oomph" in the world, and causation is just this dynamism or oomph. This talk's claims centre on a notion of meta-causation (a variety of what has been called iterated causation or higher-order causation). Meta-causation in the intended sense can also be called meta-dynamism or meta-oomph. Meta-causation is where causings—instances or portions of the world's overall causation/dynamism/oomph—are themselves causes or partial causes. That is, the very causing of B by A can itself, directly be part of what causes something C.

The distinctive claim about phenomenal consciousness is that there are good reasons to think that it and its basic "feels" or qualia are constitutively based on meta-causation. This claim is ultimately aimed at helping to solve the hard problem of consciousness. But there's no immediate reason to think that meta-causation is confined to its role in consciousness—it may have much broader significance, though this talk cannot go into what forms that significance might take.

Meta-causation or meta-oomph in the context of consciousness has already been at least implicitly raised in discussions of people possibly having direct conscious experience of causation, because that experience, if it is itself a causal matter, arguably involves meta-causation. More particularly, meta-causation is raised as an issue by recent discussions within power theory of whether we have direct conscious experience of powers at work. However, there appears to have been little discussion of the potential implication that causation, as dynamism or oomph, includes meta-dynamism or meta-oomph, and furthermore that its significance could go beyond consciousness. And as regards the link to consciousness, this talk's claim goes much further than the idea that we might be able consciously to experience causation, and makes meta-

causation constitutively foundational for consciousness itself.

The arguments for this claim start from an assumption that (occurrent, phenomenal) consciousness has a minimal core consisting of pre-reflective self-awareness (PRSA) of the sort championed by recent neo-phenomenologists. PRSA is pre-reflective in not being mediated by concepts, inference, or anything else that would ordinarily be taken to be a form of thought. But it is reflexive in that the conscious entity is reacting to itself. I argue that there are good reasons to think that this pre-reflective self-reactiveness in turn involves meta-causation.

Barone, Pamela (University of Balearic Islands), David Larrode (University of Zaragoza), Manuel Bedia (University of Zaragoza), and Antoni Gomila (University of Balearic Islands)

The minimal Turing test: reciprocal contingencies for interaction detection

Abstract. According to the second person perspective of mental attribution, the basic understanding of others' minds is direct, immediate and takes place in reciprocal interaction (Gomila, 2001, 2002; Gomila & Perez, 2017; Pérez, 2013). This basic understanding causes the experience of intersubjectivity and appears in human development before first- and third- personal attributions (Gomila, 2012; Reddy, 2008). The theory has to provide an account of the social cues interactors are sensitive to when realizing they are interacting.

The answer finds inspiration in the classical study of Murray and Trevarthen (1985) in which babies seemed to use social contingencies to distinguish when a video of their mother was online or pre-recorded. In adults, mutual recognition in online interactions has been studied with the perceptual crossing paradigm (Auvray, Lenay, & Stewart, 2009; for a review, see Auvray & Rohde, 2012). Although participants did not succeed at the task of telling apart the nature of their interactive partner, analyses of long-term dynamics revealed a particular pattern characteristic of immediate interaction suggesting that we may be able to unconsciously detect interactions (Bedia, Aguilera, Gómez, Larrode, & Seron, 2014).

We designed a new version of the perceptual crossing paradigm to match the original experiment by Murray and Trevarthen; in this way, we turned it into a minimal Turing test. Participants had to decide whether they interacted with a human or a machine –when in fact, movements of a human, either online or offline, were presented. We assume that the online version will generate the experience of interaction because it provides the cues of reciprocal contingency, while the offline will not, and so participants will treat it as the machine. As in the perceptual crossing paradigm, participants only had access to the other's movements through a sound when both agents cross (first condition). In a second condition, participants could also see how the other agent moved while, in the third condition, the evidence was again only auditory.

In study 1 three agents were employed: human online, oscillatory bot, and human offline, which was a record of a previous, successful, contingent human-human interaction (mimicking Murray and Trevarthen's offline condition). As a result: (1) participants differentiated the human agents from the oscillatory bot and correctly identified the bot as a non-person; however, (2) they judged both the online and the offline agents as persons, and (3) their performance did not improve in the audiovisual condition. Analyses of long-term dynamics revealed no specific pattern for each type of agent. Results suggest that participants based their judgments on the

pattern of movement detected (predictable, periodic trajectories or not), rather than on interactive cues. A more complex behavior of the offline agent may have induced that response.

In study 2 we only employed online and offline agents and the same three conditions (auditory, audiovisual and auditory). Results show that participants correctly identified the online agent, but only in audiovisual condition, therefore basing their response on the reciprocal contingencies of the interaction when both auditory and visual cues were provided.

Barzykowski, Krystian (Jagiellonian University) and Søren Staugaard (Aarhus University)

How intention and monitoring your thoughts influence the retrieval of autobiographical memories

Abstract: Autobiographical memories are recollections of our personal past. They may be distinguished into involuntary and voluntary memories on the basis of conscious intention to retrieve. While involuntary autobiographical memories are retrieved with no intention and effort and usually unexpectedly come to mind, voluntary memories are intended and expected. There is widespread agreement in the field of autobiographical memory that these two types of memories differ from each other in terms of retrieval processes, but the exact mechanisms are not fully understood.

In the present paper we hypothesize that involuntary retrieval depends on memories that are highly accessible (e.g., intense, unusual, recent, rehearsed), while the elaborate search that characterizes voluntary retrieval can also produce memories that are mundane, repeated or distant - memories with low accessibility. Previous research provides some evidence for this idea, which can be called a threshold hypothesis. However, in almost every prior study, participants have been instructed to report only memories while ignoring other thoughts. It is possible that such an instruction can modify the frequency and phenomenology of involuntary memories.

The present study aimed to investigate the effects of retrieval intentionality (i.e., wanting to retrieve a memory) and monitoring processes (i.e., waiting for a memory to appear) during involuntary and voluntary retrieval. We used a modified version of Schlagman and Kvavilashvili's (2008) experimental design to control the retrieval phase. Briefly, this design is a development of the word-cue method, made specifically to simulate the circumstances in which involuntary memories have been most frequently observed in daily life using naturalistic diary methods (see Berntsen, 2009, for a review). In these studies, participants report that involuntary memories usually happen while they are performing a routine task that requires little attention. In Schlagman and Kvavilashvili's (2008) design, participants were engaged in an uninteresting, undemanding vigilance task. In addition to this task, they were exposed to the presentation of short verbal phrases some of which may incidentally trigger involuntary memories. Participants in the involuntary condition were instructed to write down any spontaneously occurring memories during performance of the vigilance task. They had to interrupt the vigilance task every time they wanted to write down such a memory. In a separate voluntary condition, participants were required to intentionally retrieve a memory in response to the verbal phrases.

Across two studies, participants were instructed to (1) intentionally retrieve autobiographical memories, (2) intentionally retrieve any type of thought, (3) wait for any type of thought to spontaneously appear, or (4) wait for an autobiographical memory to spontaneously appear. This way we manipulated the monitoring processes during the involuntary and voluntary retrieval. In addition, participants also rated their mental

content both online (during the retrieval) and retrospectively (after the retrieval) on a number of phenomenological characteristics. Our results support the prediction that mostly highly accessible memories enter awareness unintended and unexpected, while memories with low accessibility rely on both intention and monitoring in order to reach awareness. We discuss the implications of the effect of intention and monitoring processes on autobiographical memory retrieval.

Baston, René (Heinrich-Heine University)

How ideological is prejudice research in social psychology?

Abstract: Haslanger (2015) claims it is not enough to change the individuals' prejudices in order to change social structures, which are the cause of systematic disadvantages for social groups. The aim of this paper is to show that Haslanger's argument has to be handled very carefully in order to avoid a fallacy. Therefore, I will firstly shed some light on the relation of prejudice and discrimination in psychology and describe Haslanger's critique of the standard narratives. Next, I will analyze a logical fallacy that some psychologists and philosophers made, that is shown by Haslanger's critique: reducing prejudices of individuals (Micro-Level) does not guarantee an overall improvement for a social group (Macro-Level). Prejudices (negative affect towards a social group) may change and negative attitudes towards social change (conservatism), can ensure that systematical disadvantages remain. Afterwards, I will show that Haslanger's argument has to be carefully spelled out in order to avoid a similar fallacy. Prejudices reduction strategies, like the procedure suggested by the contact-hypothesis, show more results than prejudice reduction alone: more empathy is triggered and cooperative behavior is promoted (Hodson et al., 2013). Taken together all effects a procedure has, Haslanger's argument can fall short if the procedure for prejudice reduction is not considered in detail.

Batisti, Filippo (Ca' Foscari University of Venice)

Escaping the isolation: From linguistic relativity to distributed cognition

Abstract: Linguistic relativity studies aim to verify whether the very condition of being a speaker of a certain natural language bears any consequence on thought or behaviour. The so-called Neo-Whorfian studies of the last 25 years or so (see Everett-2013) shared in most cases a cognitivist approach (Hunt&Agnoli-1991) in addressing the problem, with great benefit. However, more recently, this kind of approach has been put under scrutiny due to its alleged short-sightedness: Neo-Whorfian standard lab experiments may have the prerogative of involving subjects in tasks that hardly resemble the everyday practices which people actually perform. Since the lexical and grammatical items (that may influence thought and/or behaviour even offline) are learned and perpetuated during those very practices, from this point of view it may be problematic not to recreate them when testing the use of linguistic structures (Björk-2008). This kind of issues has raised concerns in ethnolinguists who stressed the shortcomings of such an approach, even without criticising its qualities or judging wrong per se (Enfield-2015). Hence a proposal for linguistic relativity studies followed, namely to consider linguistically mediated social interaction, studied with data and methods coming from Conversation Analysis, as a new locus for relativity effects (Sidnell&Enfield-2012, Enfield&Sidnell-2017). This approach, which does not address the cognitive performance of the speakers, widens the domains in which linguistic relativity effects are to be found. In other words, it would not replace the Neo-Whorfian cognitivist model. However, the heuristic import of this proposal points at the direction of a further integration, towards a new paradigm of studying linguistic relativity which would consider situations of distributed cognition between at least two subjects, conceived

as speakers-agents-cognizers, while involved in everyday linguistic practices. Such an approach, first adumbrated by Michael (2002) but never taken seriously since then, probably requires also reconsidering the extended mind model (Clark&Chalmers-1998, Clark-1998,2003,2008), starting on from two premises: first, the environment we grow up and live in is permeated by language; second, so far, cognitive science has privileged an "isolationist approach" to language, also because its obvious and legitimate methodological convenience. Such a stance analyses language post usum, thus overlooking the linguistic interactions omnipresent in everyone's life. Moreover, following Steffensen (2009), the clarkian model strays of "organism-centrism", since it conceives the extension flux only in a monodirectional fashion, from the individual's mind towards the outside, and not viceversa. This ecolinguistic paradigm (Steffensen-2011, Kravchenko-2016) aims to take seriously the social, situated and online dimensions of linguistic use, in contrast to a too much-crystalized notion of 'particular language'. This position – despite meeting the desiderata of the scholars appealing to a greater resemblance between real life situation and study settings – is not exempt from problems; not least the contraposition to the "orthodox" view in linguistics of natural languages as systems –to be substituted with notions arguably too vague, like "languaging" (Cowley-2007), "life process [vs. cognitive process]" (Boguslawska-Tafelska 2016), "sense/meaning saturation" or "existential trajectories" (Steffensen&Fill-2014). Can this paradigm serve as a good response to the questions posed by the latest trends in linguistic relativity?

Bella, Michela (Roma Tre University)

William James's "Evolutionary" Epistemology of Psychology

Abstract: In his book Darwinism and Pragmatism. William James on Evolution and Self-Transformation (2017), Lucas McGranahan offers an interesting contribution to the studies on the connections between pragmatism and Darwin's theory of evolution. He focuses on William James's psychology and philosophy to display four constitutive aspects of his reception of the theory. He argues that James's interpretation of evolution is characterized by a middle way between internalism and constructivism, making reference to the distinction used by Godfrey-Smith (1996); he underlines the fundamental role played by the mechanism of selection as a general pattern that exists in various domains and at various levels of analysis. He shows James's fallible and abductive image of science and addresses dynamic ontological continuity as a key-feature of reality that discloses real alternative possibilities. His integrated view of James's works based on evolution may be, arguably, considered new as to its completeness. However, several scholars (Schull 1992; Pawelski 2007; Franzese 2009; Pearce 2018; Klein 2016) underlined from the beginning the structural importance of Darwinism for James's perspective and the other members of the so-called "Metaphysical Club".
After a brief reconstruction of McGranahan's analysis, the second part of the paper is going to focus on the connections between Darwinism and James's pragmatist epistemology. By carefully surveying James's passage from positive psychology to his doctrine of radical empiricism, the paper proposes to show how psychology remains a lifelong field of interest for James within an evolutionary continuous ontological framework. In particular, his theoretical passage is from the continuity of consciousness to the relational continuity of the field of immediate experience. According to H. Heft and others (2017; Bella 2018), James never abandoned psychology and very soon, already in 1895, he changed his mind about the possibility to keep metaphysics (or philosophy) and epistemological issues out of psychology as a natural science. Moreover, his psychology has to be considered as a "field-oriented" conception, in so far as he was influenced from the shift of paradigm that was taking place in physics at the beginning of the 20th century. The third part of the paper makes important distinctions between James and phenomenological approaches to psychology. Despite the similarities that can be found in James's description of perception

with, for instance, existential phenomenology, and affectional facts with current interpretations of "affordances" (Chemero 2009; Heft and Richardson 2013), the Darwinian roots of pragmatism constitute an undeniable difference in this regard and help to individuate the innovative traits of this philosophical perspective.

The recent return of psychology to its roots as a life-science, integrating principles used in biology, makes Darwinism an important theory to understand the human mind and its historical development. Therefore, the pragmatist epistemology, as profoundly and originally informed by Darwinism, can furnish valid tools in the contemporary field of psychology. This does not mean identifying James's perspective with the established evolutionary psychology. The paper is rather going to point to new promising directions in which James's pragmatist epistemology can engage in fruitful discussions with other epistemologies of psychology.

Berčić, Boran (University of Rijeka)

Logical Positivists on the Self

Abstract: Logical Positivists on the Self

Logical Positivists are not known for their work on the Self. However, they had very interesting and elaborated view on the nature of the Self. The negative part of their view was primarily a critique of Descartes' Cogito, a traditional stronghold of the rationalistic philosophy. The positive part of their view was the empiricist reductionist account of the Self: the idea that the Self was a construct. Their view about the Self, as opposed to the Descartes' view, can be summed up in five points: 1. Self is not something simple, it is something composed of elements. 2. Self is not known by a direct insight, but indirectly, inferentially and gradually. 3. Self is not the Archimedean point of the knowledge, it is discovered later in the process of the rational reconstruction. 4. Self is not known a priori but a posteriori, its existence is an empirical discovery. 5. Self is not something that exists necessarily, its existence is contingent. Positivists' view was that the true nature of the Self was best revealed by its place in the constitution system, that is, by its place and role in the process of the rational reconstruction of knowledge. This is why the most interesting positivists' insights into the nature of the Self are to be found in the Rudolf Carnap's Der logische Aufbau der Welt and Hans Reichenbach's Experience and Prediction. Since they were reductionists about the Self, they had to answer the objection that is nowadays known as the objection from the unowned experience (John Campbell Past, Space, and Self). And their answer was elaborated and systematic: although the basis of our constitution system is our experience, we discover that fact later, through the constitution of the physical, autopsychological, and heterpsychological. Therefore, the reductionist analysis of the self as a construct out of experience can be carried out in a non-circular way.

Berdini, Federica (University of Bologna)

Disorders of Agency, Constitutivism, and Practical Knowledge of Intentional Action

Abstract: This paper relates cases of 'disorders of agency' to Christine Korsgaard's constitutivist theory of agency, and argues that allegedly "less than fully unified" agents (by Korsgaard's own criteria) can, in appropriate contexts, successfully exercise their agency. Although Korsgaard's work provides a rare acknowledgement that agency comes in degrees, she unduly restricts the notion of intentional action to that of rational action. As a consequence, such an account is highly demanding in terms of the development and manifestation of higher-order capacities for rational self-reflection and practical deliberation.

The disorders of agency reviewed here are associated with personality disorders, and will be used to show that while an agent's individual actions might fail to display Korsgaard's required pattern of consistency, self-reflection, and self-unification (and often, morality), her agency does not fade nor is disrupted. To account for these cases as genuine instances of agency, I suggest that Korsgaard's authorship view of agency be complemented with G.E.M. Anscombe's characterization of intentional action.

I first take up Hanna Pickard's analysis of the disorders of agency and her model of "responsibility without blame" to resist Korsgaard's assimilation of agency's constitutive norms to rational principles, and I introduce the notion of agential responsibility. On this notion, predicated upon Pickard's model, to treat someone as an agent is to treat her as accountable and answerable for her actions, which presupposes that she has (some) knowledge of what she is doing, and can exercise choice and a degree of control over her behaviour. Doing so allows me to retain important aspects of Korsgaard's account—specifically, her (practical) relational account of responsibility as answerability—while incorporating graded notions of control and conscious knowledge.

I then turn to Anscombe's characterization of intentional action in terms of practical knowledge. She describes the kind of knowledge an agent has of her own intentional action as practical, in that i) the agent's conscious awareness of her action ('agential awareness') plays a crucial role in practically orienting her throughout the execution of the corresponding action, and ii) the potential for rational control the agent has over her action puts her in a (normatively) practical relation of 'agential authority' with respect to it. Both of these features of intentional action are central to the understanding of agential capacities (knowledge, choice, control) grounding my definition of agential responsibility, and represent a significant point of continuity with Korsgaard.

Unlike Korsgaard's, Anscombe's account encourages understanding intentional actions, and actions of cognate species (such as merely voluntary actions), as structured on a spectrum of increasingly complex cognitive capacities and degrees of availability to conscious and rational control. Crucially, these capacities are developed and exercised (at any degree) through the practice of asking for and providing reasons: Following Anscombe, the description under which an agent's action counts as intentional, and therefore manifests her 'practical knowledge,' is elicited through her answers to Why?-questions. As such, incorporating elements from Anscombe's and Pickard's accounts into Korsgaard's theory provides it with the flexibility needed to better account for disorders of agency.

Błażej, Skrzypulec (Polish Academy of Sciences)

Nonclassical Mereology of Odours

Abstract: While there is a growing philosophical interest in analyzing olfactory experiences, the mereological structure of olfactorily experienced odours, i.e. odours considered in respect of how they are perceptually experienced, has not yet been extensively investigated. The paper argues that olfactorily experienced odours have a mereological structure, but it is significantly different from the mereological structure of visually experienced objects. Most importantly, in the case of the olfactory part-structure, the classical weak supplementation principle, stating that entities cannot have only one proper part, is not satisfied. This thesis is justified by referring to empirical results in olfactory science concerning the human ability to identify components in complex olfactory stimuli.

Bonalumi, Francesca (Central European University), Margherita Isella (Università Vita-Salute San Raffaele) and John Michael (University of Warwick)

Cueing Implicit Commitment

Abstract: The phenomenon of commitment is a cornerstone of human social life. Commitments make individuals' behavior predictable in the face of fluctuations in their desires and interests, thereby facilitating the planning and coordination of joint actions involving multiple agents. As such, the origin and stability of everyday social exchanges and institutions such as marriage, scientific collaboration, and employment depend upon the credibility of commitments. Despite the importance of commitment for distinctively human forms of sociality, it remains unclear how people identify, prioritize and assess their own and others' commitments. This is especially true for cases in which a sense of commitment arises in the absence of any explicit agreements. In addressing this question, Michael, Sebanz & Knoblich (2016) recently hypothesized that people's judgments and attitudes about such situations are governed by an implicit sense of commitment, which is modulated by various cues that another agent expects one to perform a particular action and may have invested effort or other costs on the basis of that expectation. In a battery of experiments conducted using M-Turk, we tested this hypothesis by presenting participants with vignettes describing everyday situations in which an implicit commitment between two agents was violated, and asked participants to judge to what extent an apology was in order, and to what extent they would be annoyed if the offending agent did not offer them an apology in such a situation. We manipulated the extent to which the offending agent had been expected to honor the implicit commitment. In Studies 1a and 1b, we presented participants with vignettes describing a scenario in which one agent had either a high level of expectation (generated by the investment of a high degree of effort into a joint activity, i.e. the High Cost Condition) or a low level of expectations (generated by a low degree of effort into a joint activity, i.e. the Low Cost Condition), and a second agent failed to remain committed. Studies 2a and 2b were designed to probe participants' normative evaluations and affective attitudes in response to scenarios in which one agent failed to remain engaged to a joint activity toward which her partner had either a high level of expectation (generated by a longer history of repeating the routine, i.e. the High Repetition Condition) or a low level of expectation (generated by a shorter history of repeating the routine, i.e. the Low Repetition Condition). In line with our predictions, the results revealed that participants made more negative normative judgments and reported more negative emotional attitudes in response to the High Cost/ High Repetition condition than the Low Cost/ Low Repetition Condition. Taken together, these results provide support of the hypothesis that people's judgments and attitudes about implicit commitments are governed by an implicit sense of commitment, which is modulated by cues to others' expectations, and by cues to the costs others have invested on the basis of those expectations.

Brozzo, Chiara (University of Tübingen & Centre for Integrative Neuroscience), Hong Yu Wong (University of Tübingen & Centre for Integrative Neuroscience), and Krisztina Orbán (University of Tübingen & Centre for Integrative Neuroscience)

How the science of action can illuminate the philosophy of action (and beyond)

Abstract: Can the science of action illuminate the philosophy of action, and, if so, how? It may be thought that the two do not (or, even, should not) interact at all.

We'll start with a theoretical defense of the possibility for the science of action to contribute to the philosophy of action, given by Hong Yu Wong. This defense will counter arguments based on the difference between the space of reasons and the space of causes.

We will then offer a number of examples of a fruitful interaction between the philosophy of action and the science of action. Philosophy of action will be interpreted in the broad sense, meaning not just answers to the question as to what distinguishes actions from mere happenings, but also other areas of philosophy that study action. One example is social cognition, which involves the study of how observed actions are understood.

We will start with the question as to whether understanding others' actions involves mindreading, i.e. the ascription of propositional attitudes. Chiara Brozzo will show how data from the science of action motivate investigating an interesting middle ground between the following two extremes: that action understanding does not involve mindreading at all and that action understanding involves the ascription of standard propositional attitudes.

We'll end with a surprising and unprecedented example of how the science of action can illuminate even areas of philosophical investigation that do not concern action at all. In particular, Krisztina Orbán will show how the notion of first person reference relies on the key ethological and developmental phenomenon of self-directed action.

Buccella, Alessandra (University of Pittsburgh)

Olympians and Vampires – talent, practice, and why "mortals" don't get it

Abstract: Cecilia Zandalasini is a 21-year-old Italian basketball player. When you watch her take a jump shot, all you can think about is that she's doing exactly what, if God existed, her assigned mission on Earth would be. Mikaela Shiffrin is a 22-year-old American alpine skier. When you watch her do a slalom run, you feel the way Cimabue must have felt watching his student Giotto draw a perfect circle, freehand.
These are only two examples of exceptional athletes about whom you can't help wondering: they most definitely have a huge amount of talent, and surely they've been working extremely hard since they were children...But is this supposed to be enough to explain what they can accomplish on the court or on the snow?
In this paper, I attempt to make sense of the relationship that there is, in the case of sports champions, between so-called "talent", i.e. natural predisposition for particular physical activities and high-pressure competition, and practice/training. In particular, I explore the possibility that our ordinary way to understand and relate such notions is wrong.
I start by trying to articulate the folk conceptions of talent and practice. Then, my argument will point out how the folk conception is over-simplistic and neglects an important aspect. Finally, I suggest a different approach in trying to make sense of how talent and practice intertwine in the very particular case of top athletes.
My approach will also highlight how, once again contrary to the folk conception, the difference between being a WNBA champion (or an Olympic gold medalist) and being an averagely skilled sports amateur is qualitative much more than quantitative. Put differently, being Cecilia Zandalasini or Mikaela Shiffrin is something an average amateur is completely cut off from, even imaginatively, as it involves what Laurie Paul (2014) calls a "transformative experience".

Bugnon, Julien (University of Fribourg)

Introspective Awareness – on what it's like to attend to one's own phenomenal states

Abstract: After decades of discredit, the topic of introspection is receiving renewed interest in recent philosophy of mind. In addition to direct discussions of the phenomenon, many contemporary debates – e.g. concerning phenomenal concepts, agency or cognitive phenomenology – seem to assume that we can attend to the phenomenal character of our own conscious states. But do we really have introspective experiences, in the distinctive sense of attending to our conscious states in such a way that the attending is itself phenomenally manifest? And if so, what is the phenomenal difference between an introspective state and its corresponding introspected state? In this paper, I explore a positive answer to the first question that relies on a specific account of the transition between having an experience and introspecting that experience, an account inspired by the notion of a Gestalt shift. The Gestalt account enables us, or so I shall suggest, to better identify the nature of introspective awareness – i.e. the kind of awareness we have of our own conscious states when we introspect – and, ultimately, to understand what introspecting consists in.

The Gestalt model provides a middle way between two prominent – yet in my view unsatisfactory – accounts of introspection. Perceptual accounts postulate an inner sense, an ability to 'turn our gaze inwards' and sense our own experiences in much the same way as we perceive external objects. However, introspective awareness seems to differ significantly from sense perception in that it does not exhibit the kind of change constancy that our perceptual awareness does (see Siewert 2012). Purely cognitive views, which construe introspection as a mere matter of forming appropriate introspective thoughts, judgments or concepts, fail for their part to account for the fact that these cognitive attitudes are based on a prior form of attention to the relevant phenomenal features they are about. On the Gestalt account, introspecting consists in shifting from a state where one is, for instance, visually aware of a shade of blue as a property of an external object, to a state where one is aware of the blue as determining the phenomenal character one's visual experience. With this picture in place, a fruitful parallel can then be drawn, I submit, with Wittgenstein's discussion of 'seeing as' in the case of visual Gestalts, paving the way to a view of introspection as a sui generis type of awareness. The resulting account is one where changes in the subject's epistemic interests inherent to her taking the introspective stance generate a cognitive transformation of the original introspected state.

Burns, Patrick (Queen's University Belfast), Teresa McCormack (Queen's University Belfast), Agnieszka Jaroslawska (University of Edinburgh) and Eugene Caruso (University of Chicago)

When is the future ahead and the past behind? The developmental emergence of time-space mapping on the sagittal axis.

Abstract: When speaking about time, individuals in Western cultures typically refer to the future as in front of them and the past as behind. Further evidence for the spatialization of time is seen in co-speech gestures and in the execution of motor-responses on temporal judgment tasks. The latter effects are sometimes taken as evidence of a deep conceptual link between the domains of time and space, beyond mere linguistic mappings. Although children appear to comprehend spatial metaphors of time from around 5-years of age it is not known whether spatial representations of time are spontaneously evoked in their everyday thinking about time as might be expected if space and time relied on shared representations. We tested children (6-to-7-year-olds and 9-to-10-year-olds), adolescents and adults on three tasks that varied in the degree to which they tapped the spontaneous versus elicited mapping of time onto space. A motor compatibility task required participants to move a slider forward or backward in response to verbally presented tensed sentences. A co-speech gesture task examined the degree to which space was utilized in

gesture to mark time, with participants asked to use gesture in explaining temporal terms. In the final task, participants made explicit time-space mappings by using a diagram of a person to indicate where past and future events lie on the sagittal plane. All three tasks gave rise to a developmental shift towards increasing spatialization of time with age. There was no evidence of spontaneous space-time mappings in the sagittal plane among 6-to-7-year-olds. However, lateral gestures congruent with a left-right mapping of earlier to later were evident at this age. From 9-to-10-years of age onwards there was evidence of consistent time-space mappings in explicit judgements and in co-speech gesture, however, only adults showed evidence of a motor compatibility effect between space and time in the sagittal plane. We discuss these results in the light of two recent claims: first, that the subjective feeling of the future as closer than the past is a consequence of representing the future as before us, and, second, that time-space mappings reflect cultural differences in the tendency to focus on different time periods.

Burns, Patrick (Queen's University Belfast), Teresa McCormack (Queen's University Belfast), Patrick O'Connor (Queen's University Belfast) and Cristina Atance (University of Ottawa)

Subjective time perception and episodic future thinking in children's delay of gratification.

Abstract: In intertemporal choice tasks participants are faced with the dilemma of trading some small immediate or short term benefit for a larger, long term reward. A preference for immediate gratification is particularly pronounced early childhood, with the shift towards more prudent decision making continuing right through to adulthood. Furthermore, individual differences in children's delay of gratification are associated with a wide range of negative outcomes in later life. Recent work with adults suggests a number of factors that may attenuate myopia in intertemporal choice tasks. In line with the hypothesis of Boyer (2008) that episodic cognition evolved for the purpose of aiding decision making, manipulations which promote episodic future thinking reduce discount rates in intertemporal choice tasks. Secondly, it has been suggested that individual differences in how far away the future feels may explain variability in the discounting of future rewards (Kim & Zauberman, 2009; 2013). These authors suggest that an exponential decay model, as opposed to a hyperbolic model, may better account for discount rates calculated over subjective time. We explored the relationship between subjective time perception, episodic future thinking and the discounting of both real and hypothetical rewards in a sample of 7-to-9-year-olds (N = 132). Steeper discounting of hypothetical rewards was strongly associated with a preference for immediate choices in a standard delay of gratification task using real rewards. Consistent with studies of adults and adolescents, children's discounting of delayed rewards was well described by a hyperbolic discount function. However, when subjective time estimates were applied, discounting was better described by an exponential decay model. There was limited variability in children's episodic future thinking and it was only weakly related to delay of gratification. We will discuss these results in relation to the claim that time perception may be a critical factor in accounting for departures from normative decision making.

Butlin, Patrick (King's College, London)

Pleasure, Pain and the Evolution of Consciousness

Abstract: According to Hedonic Interface Theory (HIT), the function of pleasure and displeasure is to provide information about value to the 'cognitive' systems responsible for goal-directed action selection. The advocates of HIT claim that the capacity for conscious hedonic experience co-evolved with goal-directed behaviour. However, this claim is in tension with evidence of conscious pain in zebrafish and

hermit crabs, since these animals seem to rely on simpler systems for behavioural control. This paper discusses this tension with reference to the broader project of understanding the evolution of consciousness.

Christensen, Wayne (University of Warwick)

Meshed control in skilled action

Abstract: One of the most striking and impressive human abilities is the capacity to respond rapidly and skilfully to complex, challenging situations. For instance, a cricket batsman adjusts to a surprising bounce and plays an unorthodox stroke that sends the ball through a gap in the fielders to the boundary. A jazz trumpeter improvises in continuous interplay with the drummer, each responding to and anticipating the other so closely that sometimes the drummer completes the trumpeter's phrases. An airliner hits a flock of birds shortly after takeoff, resulting in complete engine failure, and the pilot swiftly identifies a nearby river as the only accessible landing place in the area (a city), then with a window of only seconds flies the plane along a narrow path to a successful ditching, saving the lives of all on board. Everyday skills like cooking a meal, having a conversation, or driving a car are less spectacular but nevertheless can also involve a great deal of flexibility and context-sensitivity. These aspects of skill are not well understood. The most influential approach to skill proposes that it is largely automatic. In recent years the view that skill is largely automatic has been criticised by numerous authors. So far, however, the criticisms have been largely descriptive; that is, they have focused on arguing that cognitive control plays an important role in skilled action. There has been relatively little sustained attention to the question of why it should do so, or to the detailed understanding of how. But the systematic development of this view requires a critical examination of the rationales for skill automation employed by Fitts & Posner and Dreyfus & Dreyfus, and the construction of an alternative rationale that explains the persistence of cognitive control. Such an account would need to give an explanation of how cognitive control of skilled action that is fast, efficient and flexible might be possible. These issues will be the primary focus of the talk.

Coelho Mollo, Dimitri (Humboldt-Universität zu Berlin)

The Representational Pattern

Abstract: In this paper, I present and defend a deflationary account of representation in cognitive systems that sees representational content as capturing real patterns in nature. This underexplored approach provides new and useful tools for understanding the nature and explanatory role of representation in the cognitive sciences, opening up novel paths of enquiry that, I take, hold promise for making progress in the relatively stagnant philosophical project of naturalising representation.
(Full paper attached)

Colombo, Matteo (Tilburg University)

Miscomputation in Computational Psychiatry

Abstract: An adequate explication of miscomputation should do justice to the practices involved in the computational sciences. Unfortunately, relevant practices outside computer science have so far been overlooked. In this paper, I begin to fill this gap by distinguishing different notions of miscomputation in computational psychiatry. I argue that a satisfactory explication of miscomputation in computational psychiatry should involve a semantically laden characterisation of a computational system's interaction

with its environment. Because the mechanistic account of physical computation does not appeal to semantics, it cannot explicate a notion of miscomputation central to computational psychiatric practice.

Corrado, Maria Giovanna (University of Warwick)

Hearing Silence

Abstract: If silence is the absence of a sound, what is it to experience silence? The claim that, given that silence is the absence of a sound, silence cannot be perceptually experienced may lead some to take it that hearing silence amounts to not undergo any auditory experience. My aim is to assess whether there is an auditory perceptual experience of silence distinct from the bare absence of auditory experience. After suggesting that there is empirical support for a a sub-personal difference between the total absence of auditory experience a totally hearing impaired person supposedly undergoes and the experience of hearing silence of an average hearer when it is silent in their environment, I move on to considering whether the difference is phenomenal and, hence, whether there is a phenomenal character distinctive of hearing silence. By invoking the notion of an auditory temporal field, I will distinguish a claim that makes a minimal commitment to this notion according to which, in virtue of a temporal structural feature of perceptual experience, one hear silent pauses by hearing a sound stopping and another successively starting, and a stronger claim according to which hearing silence amounts to latching on to an environmental element, namely, an interval of time empty of sounds. I argue that hearing silence is not tantamount to the absence of auditory experience, insofar as it is possible to defend the claim that by hearing sounds temporally unfolding we can be aware of the pauses occurring between them. However, there is no sufficient argument for the stronger temporal claim to argue that it is possible to hear the absence of sounds simpliciter.

Danese, Antonio (University of Padua)

Evolutionary Psychology and Darwinian Heritage

Abstract: According to the theories that constitute evolutionary psychology, natural selection designed several independent computational modules, each of which would preside over a specific cognitive domain developed in an Environment of Evolutionary Adaptedness. On the assumption that different neural circuits have specialized during Pleistocene to solve corresponding adaptive problems, reverse engineering has been introduced to explain the evolution of the human mind and inherited social behaviors from our ancestors. Moreover, the obstinate research for an adaptive meaning for any behavioral trait has led to conceive the evolution in terms of problem solving.
In this paper I intend to show that a comparison with the Darwinian origins of the discipline reveals how evolutionary psychology misunderstands the evolutionary process and applies a simplified neo-Darwinian theory of evolution that is based on quotations of a limited number of authors (R. Dawkins, D.C. Dennett and E.O. Wilson). As a result, evolutionary psychology has focused on natural selection by underestimating the other fundamental aspect of Darwin's explanation: common descent with modifications, or the Tree of Life. Tree Thinking is the development of this Darwinian heritage carried out by modern evolutionary biology and consists in the study of genealogical kinships among species and of comparative knowledge about common descent.
I will try to show that the adaptive hypotheses of evolutionary psychology without this explanatory approach cannot be confirmed or falsified. Moreover, I will argue that efforts to understand the reasons the existence of innovative human faculties should not be directed to the research for an initial utility,

rather they should focus on the possibility of a new recombination of already existing traits.

Indeed, evolution does not contrive new traits ex novo but reuses the already available material: reusing or recycling formerly existing structures, the result of selective processes or not, is the main theoretical meaning of the modern concept of exaptation, inherited from Darwin as well. The core of my argument is that adopting within evolutionary psychology Tree Thinking and exaptation would show that not all traits set in a biological population are necessarily adaptations; the current function of a structure does not always coincide with its historical origin; reverse engineering does not allow us to elaborate falsifiable adaptive hypotheses; and, therefore, the fundamental aspect in the description of a trait is not found in its adaptive utility, in fact, some traits may prove useless and still maintain a relevant importance in providing evidence of kinship between species.

To substantiate my argument I will focus on language to show how evolutionary psychology could benefit from the explanatory pluralism of the notions of exaptation and Tree Thinking. According to evolutionary psychology, language is an optimal adaptation selected for communication. However, starting from the Darwinian thesis that the natural and genealogical classification of living species is applicable to languages and their dialects as well, I will argue that the theory of evolution can help to find a scientific explanation that conceives the evolution of language as interaction among functional adaptations and exaptations.

de Rooij, Barend Cornelis (University of Sheffield and University of Groningen)

Are stereotypes objectionable because they are inaccurate?

Abstract: A standard way of thinking about the moral wrongs associated with stereotypes traces the morally objectionable nature of stereotypes to the epistemic deficiency of being inaccurate. On this view, stereotypes inherit their moral problems from their dubious epistemic standing as generalizations about social groups that are statistically untenable. Against this influential view, I argue that stereotypes can be objectionable even when they reflect accurate generalizations about social groups. Indeed, the same grounds for believing that inaccurate generalizations about social groups are morally objectionable may give us reason to object to accurate ones. Like inaccurate generalizations, their accurate counterparts may lead us to (a) obscure the individuality of group members, (b) obscure the internal variety of social groups, and (c) perceive a real or imagined moral distance between members of stereotyped groups. Hence, the standard view fails. My argument has an important upshot for current debates in social psychological circles over whether to adopt a moral definition of 'stereotype,' capturing false generalizations about social groups, or a cognitive definition of 'stereotype,' capturing all generalizations about social groups and emphasizing the morally neutral mechanisms of generalization that underlie them. Once we notice that all generalizations can have the morally pernicious effects associated with stereotypes, the sharp distinction between morally objectionable generalizations, on the one hand, and true or accurate generalizations, on the other hand, breaks down. Accordingly, we need not restrict our attention to the set of inaccurate generalizations even if we are interested in studying the moral wrongs associated with stereotypes.

Demiddele, Kevin (Katholieke Universiteit Leuven), Tom Heyman (Katholieke Universiteit Leuven) and Walter Schaeken (Katholieke Universiteit Leuven)

Spatial reasoning from available information: where does the magic happen?

Abstract: Most research on relational order reasoning has focused on solving problems involving memorization. Also, the amount of mental models that reasoners build has been discussed a lot more than the size of these models. We have investigated the effect of mental model size in spatial reasoning

problems without memorization. In 3 initial experiments, we each time presented more than 40 adults with spatial reasoning problems, describing the spatial configuration of 5, 7 or 9 objects.

The premises remained available to reasoners when they answered the questions. Problem size was manipulated by varying, respectively, the number of irrelevant premises (Experiment 1), the spatial interval of the questioned items (Experiment 2) and a variation in problem structure that kept both number of irrelevant premises and question object distance constant (Experiment 3). All experiments had a within-subject design for the three problem size conditions.

Our hypothesis was that larger problems would lead to decreased accuracy. While this may seem a trivial hypothesis for reasoning from memory, it was not supported by the data when reasoning with available information. These simple reasoning problems did not yield a significant difference in accuracy of response, which was very high overall (up to +90%).

We performed multi-level analyses with reaction times and accuracy as dependent variables and problem size, participants and order of the problems as explanatory variables. Reaction times did show significant differences between the three conditions, in the expected direction. Throughout the 24 problems, participants developed strategies based on their familiarity with the problem structures. Looking at their first answers only gave insight into the difference in results for when such strategies are not yet applied. Our interpretation is that people reasoning from available information do not construct one single mental model, but many small models, representing only the relevant information for each reasoning step at hand.

Experiment 4 replicated the first experiment but with tasks that involved reasoning from memory, i.e. premises were no longer available when the question was presented. The results supported the claim that reasoning from memory involves a different strategy. We also collected useful pilot data for further research on hypotheses based on this interpretation.

We argue that the actual spatial reasoning, as opposed to memorizing, is situated in the intermediate model construction steps that reasoners take. We conclude that relational reasoning research in psychology has over-emphasized tasks where reasoners need to memorize the premises. We sketch perspectives on how to proceed with examining the reasoning process itself independently from memory tasks.

We believe this research is of interest to anyone interested in reasoning with mental model representations. It provides insight into which strategy reasoners use when reaching a conclusion and where the effort lies in challenging reasoning problems.

This project was accepted to the pre-registration challenge of the Center for Open Science. (https://osf.io/26hjs/?view_only=b31b2c68684540929480b093eccfbb08)

Doherty, Martin (University of East Anglia) and Josef Perner (University of Salzburg)

By Any Other Name: Theoretical Integration of Dual Naming Problems in Development

Abstract: This paper presents an account of dual naming that systematizes a diverse range of research concerning the use and conception of labels. For present purposes the focus is on developmental work on alternative naming and hypothetical word learning biases such as the much-studied mutual exclusivity bias. Our theory provides a straightforward account of the relationships between these phenomena and theory of mind development, in particular the development of false belief understanding.

We present a theoretical framework employing mental files (Murez & Recanati, 2016). This allows an integral theory of children's diverse dual naming problems and why these problems are overcome at the same time children understand false belief. A mental file is a representation whose referent is a particular object. It gathers information about that object by tracking it over time. When an object is encountered

under different guises (different appearances or verbal labels), two distinct files may be deployed for that single object. The resulting coreferential files capture different perspectives on the same object. A coreferential file can thus be used to represent people's different perspectives (e.g., belief).

Typically, the existence of different files indicates the existence of two separate objects. To mark that only a single object is involved, coreferential files need to be linked. Development of the ability to link files provides a powerful developmental explanation for success on dual labelling and perspective tasks at the same age, around 4 years: e.g., processing identity statements, overcoming mutual exclusivity (accepting different labels for an object), visual perspective taking, and understanding differences of belief.

We summarise relevant developmental research to test this claim. Children are known to unable to provide known alternative labels for objects before they can pass explicit tests of belief understanding (Doherty & Perner, 1998). However, classic word-learning findings are that even very young children generally select novel objects when they hear a novel word. This can be interpreted as demonstrating early metalinguistic knowledge or knowledge of speaker intentions, which would challenge our claim of a common mechanism developing around 4 years. In response, we present recent research showing that children's referent selection ability is initially inflexible: for example, young children select a novel inedible object (a plastic gadget) even when given pragmatic information that a familiar edible object (an apple) is the referent ("I'm hungry, give me the zav"; Gollek & Doherty, 2016). Similarly, when taught two novel names (by different speakers) for a novel object, young children will only respond to one of them in a particular situation, whichever happens to be used first. Nevertheless, they will later correctly use the refused word when there is no competing label (Karadaki & Doherty, 2018). Both tendencies are overcome once children able to pass the false belief task.

Our conclusion is that mental files provide the best current framework for understanding referent selection and mutual exclusivity in children. and its relation to conceptual pacts and similar phenomena in adults.

Doulatova, Maria (Washington University in St. Louis)

Neuroscience of Mindreading: A Case for a Plurality of Specialized Mindreading Mechanisms

Abstract: Mindreading--the capacity to explain others' behavior and mental states in terms of their mental states--has been at the forefront of psychological and philosophical debates alike. Specifically, (i) is mindreading achieved by a type of theoretical interpretive method (TT) or does it involve simulation (ST) akin to putting yourself in another's shoes? (ii) which cognitive mechanism(s) subserves these methods? (iii) is a single method and mechanism used for all types of mindreading?
Carruthers (2011) offers one answer to the above questions: a single sub-personal interpretive or theory-based method, subserved by a single cognitive mechanism, is used for all types of mindreading (i.e. Interpretive Sensory Access Theory,(ISA)
I offer two empirical arguments against Carruthers' one-for-all hypothesis.
If the one-for-all hypothesis obtained, we would not see double dissociations - of the kind indicative of distinct cognitive mechanisms - in neural activations during two mindreading tasks (i.e. mindreading similar others versus mindreading dissimilar others). Which mechanism is engaged seems to vary as a function of the mindreader's perceived closeness to the target. While perceived closeness should be irrelevant to TT and ISA, it could be relevant ST. Given that ST posits that we use our own cognitive capacities to simulate what it would be like to be in the target's shoes, it would be easiest to do if the target is not much different from the mindreader. Thus, contra Carruthers, ST plays an important role in mindreading.
Furthermore, an unspecialized mechanism could be disastrous given the contrary demands of socializing

with ingroup versus outgroup members. If the intimacy-promoting mindreading mechanism reserved for mindreading ingroup members were instead used for mindreading outgroup members, aggressive actions necessary for ingroup/outgroup competition would be irrevocably compromised .

The goals of social interaction with ingroup members are radically different from those of interactions with outgroup members. This suggests that predicting and explaining behavior and mental states are not the only outcomes of mindreading. Ingroup interaction should promote bonding, whereas outgroup interaction should promote the opposite. Besides predicting and explaining behavior, mindreading also fosters intimacy. On the other hand, in mindreading your competitor, you're best served to maintain psychological distance. Since interactions between members of different groups could be costly to both parties, and ingroup/outgroup competition often calls for calculated actions, mindreading dissimilar others should not promote bonding.

Given what we know about TT/ISA versus ST, it is arguable that the goals of fostering intimacy are better served by the latter while the goals of promoting competition by the former.

This is exactly what we see. Being prompted to use ST in situations calling for calculated or utilitarian actions results in sub-optimal outcomes. If the target's mental states are inferred using TT and not ST, these mental states are better "quarantined" in the mindreader's mind.

Inference to the best explanation suggests that there are at least two specialized mindreading mechanisms, each dedicated to particular mindreading tasks depending on (i) mindreader's perceived closeness to the target, and (ii) working memory demands prior to and following the mindreading task.

Dranseika, Vilius (Vilnius University)

Quasi-memories and artificial memories are memories

Abstract: In this paper, I present new data bearing on two constraints that are often taken to be essential features of our ordinary use of verb "remember": strong previous awareness condition and the factivity constraint. Study participants agreed that the agent in the experimental vignette remembers "other people's memories", which suggests that ordinary notion of memory is not bound by the strong previous awareness condition. Concerning factivity constraint, both artificial memory (event that did not occur) and quasi-memory (event that did occur, but to someone else) vignettes involve violations of factivity. Nevertheless, in both cases study participants agreed that the agent remembers. In summary, the data provided in this brief report provide some evidence that allows to motivate skepticism concerning whether the strong previous awareness condition and the factivity constraint are essential features of our ordinary use of verb "remember". To the best of my knowledge, this is the first attempt to address empirically the question whether these two constraints are features of our ordina

Ebel, Sonja J. (Max Planck Institute for Evolutionary Anthropology & University of St Andrews), Christoph J. Völter (Max Planck Institute for Evolutionary Anthropology & University of St Andrews) and Josep Call (Max Planck Institute for Evolutionary Anthropology & University of St Andrews)

Functional fixedness in great apes

Abstract: We usually assume that experience enhances problem-solving performance. However, a change in the structure of the problem may require overcoming past experience. For example, humans struggle to use objects in unfamiliar functional contexts since they habitually use them for specific purposes only

("functional fixedness effect"). Little is known whether our closest living relatives, the nonhuman great apes, are vulnerable to this effect as well. We therefore investigated if bonobos (Pan paniscus), chimpanzees (Pan troglodytes) and Sumatran orang-utans (Pongo abelii) would exhibit a functional fixedness effect. In experiment 1, apes (N = 23) either received prior experience with the brush-end of a tool (i.e., dipping juice from a container) or not. In the test, apes were presented with the same container but this time the dipping option was blocked. Instead, they had to use the pointy end of the same tool to puncture a hole in the container to access the juice. In experiment 2, apes (N = 24) either received prior experience with a hose (i.e., drinking juice from a container) or not. In the test, apes were presented with a horizontal tube with blockages close to both openings that required selecting the flexible hose to poke out the food reward. In experiment 3, apes (N = 35) were first fed with bread sticks, pieces of bread sticks or no bread sticks at all. In the test, they had to use the bread stick to rake-in grapes. Results indicated that prior experience shaped apes' manipulation style (experiment 1) and that individuals with prior experience performed worse in the test than apes without such experience with regard to success and latency (experiment 2). Moreover, apes who had not experienced the bread stick as a food item before were more likely to use it as a tool (experiment 3). In conclusion, our findings indicate that nonhuman great apes, like humans, are vulnerable to the functional fixedness effect. We also discuss implications for apes' understanding of tool functions.

Falck, Andreas (École Normale Supérieure and Lund University), Tiziana Zalla[1] (École Normale Supérieure), Ghislaine Labouret (Université Paris Descartes), Manali Draperi (Université Sorbonne Paris Cité) and Brent Strickland (École Normale Supérieure)

Our own perceptual experience, but not that of others, influences object detection

Abstract: Do we automatically compute others' false beliefs? Kovács, Téglàs and Endress (2010) provided evidence that indeed we do. In their experiments, infants and adults appeared susceptible to interference from an animated agent's belief about the location of an object: participants' responses in detecting the object were faster when the agent's belief was consistent with the object's location, compared to when it was not. Subsequent experimental work however has called the automaticity of belief representation into question. In particular, it has been claimed that Kovacs's et al. original result was based on a confound in the task and is thus not informative about processes of belief representation (Phillips, Ong et al. 2015). However, the purported failure of a single experiment to find evidence a process does not imply that it does not exist, hence the present study.

We present results from two online experiments in which participants were asked to quickly indicate the location of an object just after it appeared from behind one of two occluding screens. The object's final location could be congruent or incongruent with both their own belief and the apparent belief of an animated agent who witnessed part of each episode. In each experiment we manipulated the predictive potential of the agent's belief (between subjects): the agent's belief could correspond to reality in either 25%, 50% or 75% of trials, depending on condition. This design allowed us to investigate whether participants would learn to represent others' beliefs when statistically relevant to the task.

The first experiment (N = 307), revealed clear effects for participant belief only. When the outcome was consistent with participants' beliefs, participants were faster and less error prone. However, the results revealed no such facilitation of reaction speed and accuracy by the agent's belief in any of the "predictive" conditions.

---

[1] Very sadly, Tiziana Zalla passed away on 28th April 2018.

A second experiment (N = 208) included a number of "catch trials" in which participants had to indicate where the agent had last left the object, which was where the agent would believe the object to be. This task allowed us to ask whether the failure to find an effect of the agent's belief in the first experiment was due to an absence of automatic belief computation, or due to an insensitive experimental design. The results here revealed an effect in the expected direction for error rates (though not reaction times) in the 'predictive' condition: participants made fewer errors if the agent's belief was true compared to false, but only when the agent's belief predicted the object's location. These "belief-like" results are best interpreted as a bi-product of participants preparing a response for catch trials. More generally, they suggest that our experimental paradigm in Experiment 1 is sensitive enough to find evidence of automatic belief representation if such a process exists.

Overall, these results suggest that rapid judgments about object locations are not automatically affected by other people's apparent beliefs, but that similar effects can be elicited by directing participants' attention to task-related stimuli.

Farina, Mirko (King's College, London)

Taxonomising phenotypic plasticity: the crucial role of cultural plasticity in the evolution and development of human cognition

Abstract: ESPP 2018

Abstract
Pigliucci defines phenotypic plasticity as: 'the property of a given genotype to produce different phenotypes in response to distinct environmental conditions' (Pigliucci 2001 p.1). A similar understanding is endorsed by Marie Jane West-Eberhard, who defines phenotypic plasticity as the capacity 'of a single genotype to produce more than one alternative form of morphology, physiological state and/or behaviour in response to environmental conditions' (West-Eberhard 1989, p. 249). Phenotypic plasticity is thus generally understood as the sum of mechanisms and causes determining the ontogenetic changes that take place in an individual throughout the lifespan.
Building and expanding on this research, in this paper I offer a set of preliminary steps towards the development of a multidimensional taxonomy, which highlights and foregrounds: (a). the different types or varieties of plasticity that characterise phenotypic plasticity and (b). the contribution of different forms or varieties of phenotypic plasticity for the development and evolution of human cognitive behaviour.
To achieve these goals, I begin (section 2) by discussing one of the most researched kinds of phenotypic plasticity in the literature; the notion of neural plasticity. Neural plasticity can be defined as the capacity of the brain to reorganize its neural connections in response to changes in behaviour and environment. I demonstrate the crucial role accorded to neural plasticity in modern neuroscience by analysing two case studies (involving taxi driving and sensory substitution devices). Having done this, I go on to show the limits of an exclusive focus on neural plasticity for a proper understanding of human evolution and cognitive development. I claim that there are many other important kinds of phenotypic plasticity (besides neural plasticity) that are relevant to us humans and argue that these types of plasticity require careful systematization.
In section 3, I define and taxonomise six major types of phenotypic plasticity:
immunological, morphological, physiological, neural, cognitive, and cultural and offer a preliminary way to group them together. I then organise these different types of plasticity along a different dimension - cognitive status (3.2). This is supposed to highlight the contributions of those six varieties of phenotypic plasticity to understanding human cognition. In section 4, I focus on the three varieties of phenotypic

plasticity (neural, cognitive and cultural) that my previous analysis (3.2) has revealed as being of primary importance for understanding human cognition and reflect on their similarities and differences as well as the interrelations between them. I then show how phenotypic outcomes are transmitted (via genetic and non-genetic inheritance) across generations and argue that cultural plasticity is a crucial non-genetic, developmental channel for the trans-generational transmission of such outcomes. In section 5, I discuss two case studies, involving arithmetical cognition and underwater vision that may be taken to highlight the fundamental role of cultural plasticity in human cognition thus providing empirical support for the argument I sketched in section 4. I conclude (section 6) with a brief summary of what I have achieved.

Fazekas, Peter (University of Antwerp), Georgina Nemeth (Eotvos University) and Morten Overgaard (Aarhus University)

White dreams are made of colours: What studying contentless dreams can teach about the neural basis of dreaming and conscious experiences

Abstract: White dreams, also called contentless dreams, are reported dream experiences without the ability to recall any specific content. In serial awakening paradigms, approximately 30% of recorded answers report white dreams. White dreams are typically interpreted as forgotten dreams and are sidelined as not being particularly informative with regard to the nature of dreaming. In this paper, we propose a paradigm shift with respect to the status of white dreams arguing that focusing on this phenomenon can reveal fundamental insights about the neural processes that occur in the dreaming brain.

Recent studies comparing the EEG correlates of dreamless sleep, white dreams, and dream experiences with reported content argue that the neural correlates of specific dream contents (both in REM and NREM sleep) are local increases in high-frequency EEG activity in a so-called posterior hot zone (present only in the case of dreams with reported content), while a decrease in low-frequency EEG activity in the same posterior hot zone is the marker of the occurrence of experiences themselves (characteristic of both white dreams and dreams with reported content).

In this paper, we argue that these claims are dubious and are not supported by the data presented in the original studies.

First, we point out that, on the basis of what is known about neuronal down states, decreased low-frequency activity is not a good candidate for being the neural correlate of consciousness, neither can it be a reliable indicator of the occurrence of dream experiences.

Second, we argue that analysing a tension in the data presented (that has been overlooked in the original studies) reveals that high-frequency EEG activity in the posterior hot zone is increased in the case of white dreams as well, only not to the same extent as in the case of dreams with reported content. Increased high-frequency activity indicates higher firing rates, and thus higher intensity populations codes.

Third, we argue that the intensity of the neural population codes underlying conscious experiences determines the saliency of the corresponding conscious experience: lower intensity results in less clear, less salient conscious experience. Consequently, white dreams are less vivid, lower quality experiences of specific dream contents. Less vivid, lower quality experiences are harder to recall, resulting in similar

reports, regardless whether they occur in sleep or during wakefulness.

This paradigm shift provides a new framework for thinking about white dreams, NREM mentations, REM dreams and wakeful experiences as belonging to the same continuum. It also redefines the connection between dream recall and dream production, and reframes the classical debate concerning the role of REM and NREM sleep in dreaming in terms of the constraints that the physiological characteristics of REM and NREM stages of sleep impose on the maximal intensity of relevant neural activity patterns. We propose that combining multivariate pattern analyses techniques with the serial awakening paradigm in future research could shed further light on the relation between the intensity of cortical activity and the quality of dream experiences.

Ferretti, Gabriele (University of Florence) and Brian Glenney (Norwich University)

Pictures, Holes, and Shadows: What Would Molyneux Do?

Abstract: Pictures, Holes, and Shadows: What Would Molyneux Do?
New concepts, distinctions, and technologies have led to numerous reformulations of Molyneux's question, whether people with blindness can immediately re-identify cube and sphere shapes previously familiar to touch if given 'new sight' (Glenney 2013; Jacomuzzi et al. 2003). We consider three new reformulations based on recent work in the philosophy of perception of pictures, holes, and shadows, in response to the claim that newly sighted subjects cannot, at first sight, rely on robust shapes representations (Schwenkler 2013) required for identifying real-world shapes that undergo change with movement. Call this problem the 'robust shape perception problem,' henceforth RSP. While others have presented a similar bestiary of Molyneux question reformations (most recently Matthen and Cohen (draft)), our proposal offers two advantages: (1) a new perspective on perception of pictures, holes and shadows and (2) new possible experimental paradigms for answering Molyneux's question.

Pictures. To avoid RSP, Schwenkler (2013) proposes presenting subjects with pictures of shape rather than real objects. Criticism of this proposal (Connolly 2015; Cheng 2015) invites reformulating Molyneux's question in the context of picture perception. The current debate on picture perception suggests that the visual system represents two things: the surface of the picture and the object represented in the picture (Nanay 2017; Lopes 2005; Hopkins 2012). Can this twofoldness be identified at first sight? In other words, would a newly sighted subject identify a real cube from a pictorially depicted cube?

Holes. Hole perception (Casati 2006, 2007) requires distinguishing what in vision science is called a palpable material object from immersive negative space—a hole (for these two notions see Vishwanath 2011, 2014). Can a blind subject, able to tactilely distinguish a coin from a hole of the same dimensions, identify the hole if given new sight? Understanding the extent to which immersive negative space and palpable space can be given to both the sensory modalities can show the extent to which subjects are in the position to discriminate shapes in the context of RSP. Their success supports the view that "vision is palpation with the look", originally suggested by Merleau-Ponty (1968) and now vision science (O'Regan and Noë 2001).

Shadows. Though RSP might be avoided by using pictures, or by subject or object movement to reveal otherwise hidden faces of the object (Schwenkler 2013), we know that shadows cast by real objects provide important visual cues aiding shape identification (Casati 2007; Resnik et al. 2004). Consider a blind subject who has knowledge of basic optics and luminance distribution. Might this subject, upon new sight, detect the occluded luminance from the object being illuminated, and thus identify the cube's flat trapezoidal

shadow? If so, this suggests that the solidity of the object can be distinguished perceptually, or at least inferred by its ability to occlude light, rather than an ensemble of different cube faces within the same spatial location (Fine et al. 2003), illuminating how the visual system recognizes shapes by distinguishing objects from their shadows.

Ferretti, Gabriele (University of Florence) and Silvano Zipoli Caiani (University of Florence)

Motor Mediators

Abstract: Motor Mediators
One of the most interesting issues in contemporary philosophy of the mind is determining how intentions relate to motor representations (henceforth: MRs). Although both intentions and MRs are mental components of action, the former are usually conceived as having a propositional format (e.g. Bratman 1999; Mele 1992), whereas the latter are built in a motor format (Jeannerod 2006; Butterfill and Sinigaglia 2014; Ferretti 2016; Zipoli Caiani and Ferretti 2017). Standard accounts assume that a full explanation of the purposiveness of action requires coordination between intentions and MRs so that it is possible to proceed from the representation of the former to the representation of the latter (Bach, 1978; Searle, 1983, Pacherie, 2000).
However, the philosophical enterprise of offering such an account presents several difficulties: since intentions and MRs have different formats, it is not clear how we can explain the way in which they interlock. Intentions can be featured as premises or consequences in practical reasoning by virtue of their format. MRs are built in a motor format. So, how can these two processes interlock? This problem is known as the interface problem and a possible solution to it has been recently addressed by Butterfill and Sinigaglia (2014) (henceforth: B&S) as well as by Mylopoulos and Pacherie (2016) (henceforth: M&P).
In this paper, we propose a new account concerning the interlock between intentions and MRs, showing that the interface problem is not as deep as previously proposed. Before discussing our view, we report the ideas developed in the literature by those who have tried to solve this puzzle before us. The paper proceeds as follows. First, we respectively address the proposals by B&S and M&P and argue that both solutions entail a translation between representational formats, which, however, both accounts aim to avoid. Then, we present our brand-new claim, supported by empirical evidence, according to which intentions and MRs partially share the same motor format, inasmuch as executable action concepts are naturally represented in the agent's motor system together with action's outcomes. Indeed, since intentions are constituted by executable action concepts and since there is empirical evidence that action concepts are represented (and, thus, built) in the same motor format as action outcomes, the interlock between intentions and MRs no longer constitutes a problem: executable action concepts are the motor mediators between intentions and motor representations.

Fischer, Eugen (University of East Anglia) and Paul Engelhardt (University of East Anglia)

Salience bias in polysemy comprehension: Psycholinguistic phenomenon and philosophical relevance

Abstract: We present experimental evidence of a previously unrecognised cognitive bias and argue for its philosophical relevance. The 'salience bias' arises where polysemous words have a dominant sense more

salient than all others, and the stereotype associated with the dominant sense is functional for the interpretation of less salient uses. The bias leads to contextually inappropriate stereotypical inferences, which competent speakers go along with, even when they know that the inferences at issue are inappropriate. The talk presents a psycholinguistic study that documents inappropriate stereotypical inferences and deploys the results in the spirit of experimental philosophy's 'Warrant Project' (Stich & Tobia, 2016) to assess spontaneous judgments (intuitions) about verbally described cases.

Polysemous verbs activate most rapidly and strongly the stereotypes (situation schemas) associated with their most salient sense (Fein et al., 2015; Giora, 2003), where 'salience' is a function of exposure frequency, modulated by prototypicality. Less salient uses are often interpreted by retaining that dominant schema and suppressing contextually inappropriate components ('Retention Strategy') (Giora et al., 2014). But frequently co-occurring components of a situation schema activate others (Hare et al., 2009; McRae et al., 2005). Therefore irrelevant schema components cannot be completely suppressed while retaining relevant components. Contextually irrelevant components that are only partially suppressed then support contextually inappropriate inferences. We hypothesise that competent language users will go along with such inferences, despite knowing they are inappropriate.

We examine this hypothesis by studying inferences from the verb 'to see', whose visual sense is far more salient than all others (Fischer & Engelhardt, 2017). This talk presents an experiment that combined reading-time measurements with plausibility ratings to (1) garner evidence that the Retention Strategy is used to interpret less salient uses of the verb, and (2) document contextually inappropriate stereotypical inferences from such uses. In a pre-study, participants assessed the propriety of spatial inferences (from 'S sees X' to 'X is in front of S', licensed by the dominant visual sense) from various less salient uses. Spatial inferences from epistemic uses ('I see your point') were deemed least appropriate, and strongly inappropriate. Our experiment used a cancellation paradigm (Harmon-Vukic et al., 2009) to examine spatial inferences from epistemic uses of 'see'.

In our sentences, we manipulated verb, object (concrete/visual vs abstract/epistemic) and post-object context (placing the object in front of or behind the agent). We measured first-pass, second-pass, and total reading times for verb, object and context regions, and elicited plausibility ratings for items. Reading times for verbs and objects of 'see'-sentences were consistent with use of the Retention Strategy. Reading times for contexts were consistent with the hypothesis that participants made spatial inferences from both visual and epistemic uses of 'see', despite the acknowledged impropriety in the latter case. Plausibility ratings confirmed that inappropriate inferences were insufficiently suppressed and affected further cognition. These findings empirically support a fresh version of the metaphilosophical 'esotericity thesis' (Weinberg, 2015; Williamson, 2007): Intuitions about unusual (stereotype-divergent) cases are unreliable, when these cases are described with familiar words (possessing dominant senses), given special (less salient) uses.

Fuller, Timothy (Yonsei University)

Erasing Race? Models of Racial Conceptual Revision

Abstract: Race eliminativists advocate for a major conceptual transformation. Several eliminativists have compared their proposed conceptual revision to those involving phlogiston, caloric, humors and witches. For such concepts and accompanying beliefs, an erasure model is plausible—the relevant concepts and beliefs were likely removed if not over an individual's lifespan, at least over generations. But conceptual revision often takes a different form. Rejected concepts and accompanying beliefs often co-exist alongside

more newly acquired concepts and beliefs, both over an individual's lifespan and over generations. On this model, more newly acquired concepts and beliefs cognitively suppress, though not invariably, the influence on behavior that rejected concepts and beliefs might exert. A co-existence model likely characterizes conceptual revisions involving impetus theories of motion, final causes and animacy-based criteria for living things.

Which model of conceptual revision would be more likely to characterize the rejection of race? Evidence concerning race encoding appears to support an erasure model, and has been appealed to by some race eliminativists. However, evidence concerning the formation and subsequent invariance of implicit biases instead supports a co-existence model. In light of the total available evidence, I argue that racial conceptual revision would most likely be characterized by a co-existence rather than erasure model. Even if society's members were to consciously reject race, racial biases would likely nevertheless persist at the implicit level both over individual lifespans and over generations. Further, I argue that this undermines a prominent argument for eschewing racial categories which maintains that rejecting race is a necessary and desirable step towards achieving a racially just and egalitarian society.

Why view research on implicit biases as more compelling than research on race encoding? In short, because extant studies on racial encoding and recall lack ecological validity. In these studies, racial encoding and recall is undermined when physical appearance markers for racial identification are de-linked from "coalitional alliances" with subjects. But research on implicit bias formation suggests that such conditions rarely obtain. Rather, patterns of implicit ingroup and outgroup bias form and vary across societies because racial appearance markers are linked to disparate levels of wealth, income and social status. Further, research on implicit biases suggests they are invariant over the lifespan and causally implicated in some of the very racial disparities that cause their formation. Thus, in order for racial categories to not be encoded at the implicit level, wealth, income, and social status would have to be divorced, in children's social environments, from racial appearance markers. But this would be a world in which society has already achieved most of its goals in terms of racial justice and egalitarianism, not a world that requires the elimination of racial categories to achieve those goals. Moreover, research on implicit bias interventions suggests that the cognitive preservation of racial concepts at the explicit level may well be required for mitigating implicit discrimination. Thus, the most plausible model of racial conceptual revision implies that we have equal or greater moral grounds for conserving as opposed to rejecting race.

Gärdenfors, Peter (Lund University)

Epigenetic of semantic domains

Abstract: Children learn a language without effort. A fundamental question is how children know what to learn. When a new word is uttered in a situation, there is often a multitude of objects, features of objects and features of the on-going events that could be the meaning of the word. How does the child select the right meaning?
This paper argues that early language learning in children emerges from five primary semantic domains: Space, objects, actions, number and events. These are called 'core knowledge domains' by Spelke and Carey, but unlike their nativist view, I take an epigenetic position and claim that the domains are learned. At an early age children learn to transform a complex sensory input to concepts and categories. My central thesis is that the conceptual structures that emerge in children's cognitive development are based on detecting invariants in the sensory input. To some extent, my analysis follows the program of Gibson,

although my approach is more cognitively focussed.

The invariant structures involve a reduction in dimensionality of the sensory information. I argue that space, object, action and number domains are very natural outcomes of a reduction of sensory information in terms of invariants. Each of these semantic domains corresponds to a separate set of invariants. For example, our representation of space is characterized by the fact that distances between locations remain invariant. The cognitively constructed space is also invariant of eye, head and body orientation and of body location. In contrast, objects are can have varying locations in space, but are determined by shape invariants, among other things. The brain learns to pick up the invariants for space via processes in the dorsal "where" path and the invariants for objects via the ventral "what" path. Actions are determined by invariants of force patterns, which entails that who performs the action or where it is done is irrelevant. The number of a collection is invariant of the locations and identities of the objects in the collection. Finally, events combine information about agent, patient, action and result that contain information from the other domains.

The semantic domains are naturally modelled by conceptual spaces that are based on similarity so that the more similar two entities are, the closer they are located in the conceptual space. Categories can be described as regions of such spaces. The categories then form the basis for the meanings of early words that children learn. I argue the semantic domains can be used to generate an epigenetic model of the acquisition of word classes. Categories based on the object domain are typically expressed by nouns; categories based on the action domain are expressed by verbs; and relational categories based on the space domain are expressed by prepositions (although many languages use other means to express spatial relations); and categories based on the number domain are expressed by quantifiers and numbers. In this way a semantic foundation for the most common word classes is established. Finally, the event domain forms the meanings of declarative sentences.

Giananti, Andrea (University of Fribourg)

I know how I know: perceptual experience, self-awareness and self-knowledge

Abstract: When a subject has perceptually grounded knowledge, she typically knows how she knows what she knows, and is able to revert to perceived items in reason-giving practices. What explains this ability? In this paper I focus on visual perceptual experience, and I argue that paradigmatic cases of visual perceptual knowledge are such that, when a subject knows that p by seeing that p, she knows that she sees that p; I also argue that the truth of this thesis is grounded in phenomenological facts concerning awareness and self-awareness.

Giustina, Anna (Institut Jean Nicod)

Introspection without Judgment

Abstract: The topic of this paper is introspection of phenomenal states, i.e., the distinctively first-personal method through which one can form beliefs about the phenomenology of one's current conscious mental states. All and only the subject's current conscious states with phenomenology (i.e. states there is something it is like to be in) are potential targets of introspection of phenomenal states.
Introspecting is something we can actively do, a mental act we can perform. Much of the literature on introspection seems to assume that the product of such an introspective act is always an introspective judgment. Accordingly, the debate revolves around the nature (origin and content) and epistemic features (reliability, justification, possession of special epistemic properties such as infallibility and immediacy) of

introspective judgments. Here I will argue that there exists a pre-doxastic product of introspective acts, a mental state without a propositional structure and a conceptual content more psychologically fundamental than introspective judgments. More precisely, I will argue that two different kinds of introspective state should be distinguished: one which involves recognizing and classifying the introspected phenomenal state as an instance of a certain experience type, and another which does not involve such classification. Whereas the former is potentially judgment-like, the latter is not. I call them, respectively, 'reflective introspection' and 'primitive introspection'.

The purpose of this paper is to argue that primitive introspection is a psychologically real phenomenon. I first introduce the distinction and provide some preliminary motivation to accept it. Then, I present my central argument for the existence of a non-classificatory kind of introspective state, what I call the 'argument from phenomenal-concept acquisition'. I argue that, with some caveats, the classificatory/non-classificatory distinction maps onto the conceptual/non-conceptual distinction, where an introspective state is conceptual iff it depends on the deployment of some phenomenal concept(s) already possessed by the subject prior to introspecting, nonconceptual iff it does not depend on the deployment of any phenomenal concept possessed by the subject prior to introspecting. Then, I argue that a view according to which all introspective states are conceptual in this sense entails an implausible form of nativism about phenomenal concepts. Roughly, the idea is the following. It is natural to think that we acquire phenomenal concepts by introspection. Now, if one's introspective state is conceptual, then one can only have it if one already possesses the concept(s) associated with what is introspected. If all introspective states are conceptual, then the possession of all phenomenal concepts must be prior to the possibility of introspecting. But if this is the case, then phenomenal-concept acquisition is mysterious. The full argument is much more complex, however, and the bulk of the paper is devoted to developing it thoroughly.

Glenney, Brian (Norwich University)

Are there two moons in a moon tube? Sizing up the moon illusion

Abstract: The horizon moon appears larger than the elevated moon, yet the proportion of the visual field it takes up is the same. This phenomenon is the moon illusion. The prevailing explanation, though itself problematic, is that distance cues from intervening terrain unconsciously increase perceived distance of the moon, which in turn increase the perceived size experience, which in turn decrease perceived distance experience, such that the moon appears larger/nearer when intervening objects are also seen (Kaufman and Rock, 2007). I present preliminary data from a moon tube, a device that allows subjects to measure the perceived size of the moon, that suggests "intervening terrain" is not necessary to produce the moon illusion.

Berkeley claimed that the moon illusion persists when we block out intervening objects through a tube, (NTV 1709: 70) and Molyneux, Euler, and half a dozen others agree. Others including Reid, Malebranche, Porterfield, and another half dozen claim that moon-tube-looking extinguishes the illusion, (Plug and Ross 1989: 19; Ross and Plug 2002: 120). Call this tube-gate. Is it not resolvable empirically?

With one eye looking through an empty telescope tube, subjects adjusted an aperature fitted to the tube's end to make it exactly surround the perceived-size of the moon. In two separate tests in varying conditions, the measurements of the aperature for the horizon moon were generally larger than of the zenith moon (See Figures 1 and 2) suggesting that terrain is not a necessary condition for the moon illusion, a claim that would disconfirm the prevailing account. However, other measurements suggest that the moon illusion was erased in others, suggesting the empirical validation of tube-gate What might explain the illusion's persistence in some and extinction in others?

I conclude by suggesting that the horizon and zenith moon may be two distinct objects of perception with distinct size expectations. The two moon view finds initial support by recognizing that that it is unlikely that the percieved size of the moon would noticeably change if its transit were viewed unceasingly (Simons 2005) and that the size differences of the horizon and zenith moon are not apparent in photographs (Enright 1975). In addition, the role of normative or canonical size generates anticipations that are crucial for the perceived size of objects (Konkel 2008). The two moon view is suffiently agile for explaining tube-gate, allowing some subjects to be measuring one moon and thereby having measurements that show the illusion to be extinguished and others two moons that explain the illusions persistence.

Goffin, Kris (University of Antwerp and Ghent University)

Better Scared Than Sorry: A Pragmatic Account of Emotional Representation

Abstract: Emotional representations seem to be very unreliable. For instance, we are often afraid when there is no danger present. This is the emotional unreliability problem: if emotional representations are so unreliable, what function do they have in our representational system? I present the pragmatic account of emotional representation, which shows that emotional representations are reliable in the "pragmatic sense". Instead of seeking an optimal balance between minimizing inaccurate representations and maximizing accurate representations, emotional representations are only reliable in so far as they maximize accurate representations. They detect phenomena, based on little evidence, which implies that they often present false alarms. When it matters, however, an emotional representation will detect the phenomenon. Careful reasoning aimed at minimizing false beliefs results in fewer false alarms, but it might miss some phenomena that emotional representations would have detected. Often one is better scared (and possibly wrong) than sorry.

Grčki, David (University of Rijeka)

Making decisions over time: can a preference change be rational?

Abstract: One of the key problems in the philosophy of rationality is how we make decisions. We make decisions every day. From mundane decisions about what to drink on our lunch break to more important decisions which have further consequences on our lives. In order to analyse our decision making in the context of the philosophy of rationality we certainly need two things (two conditions): (i) normative assessment and (ii) explanation and prediction of behaviour. The normative dimension of rationality is prescriptive, i.e. it tells us what people should do while the descriptive dimension of rationality deals with what people are actually doing. In other words, it tries to explain and predict human behaviour.
Here, I am interested in a specific case of decision making problem: dynamic choice. Dynamic choice is a type of decision problem in which one's choice/choices are spread over time. People engage in sequences of choices that are not always reducible to a series of independent, individual choices. They make choices about how they will choose, and they make choices in the light of earlier commitments to choose in certain ways. They make plans for the future and they have a degree of concern for the plans that they have made in the past. This is the notion of sequential choice.
As the normative dimension of rationality I take the abstract axiomatic model of decision making (von Neumann and Morgenstern 1944) and as the descriptive dimension of rationality I take several commonsensical examples to illustrate real decision making in practice.
I am trying to find a reasonable middle ground between normative and descriptive in order to solve the problem of dynamic choice. Intuitions tell us (at least most of us) that some preference changes should be

considered as rational while others should certainly be considered as irrational. Rational preference changes are those associated with the agent changing his mind in situations that do not seem problematic, like wanting coffee at one moment and choosing tea the moment after. Preference changes that should be considered irrational are those carried out arbitrarily and repeatedly by an agent. The question I try to answer is what are good criteria for the distinction between rational and irrational preference changes, that is, is there an appropriate theoretical ground that could justify our intuitions in support of such a distinction. It has been shown that formal decision theory cannot help us. Choices can always be presented is a way that breaks the independence axiom of formal decision theory leading into sequential inconsistency (Bermúdez 2009).

I propose differentiating between rational and irrational preference change based on reasons. I use Mercier and Sperber's (2017) model of reasons as social constructs to distinguish rational from irrational behaviour. Mercier and Sperber's model together with their concepts of intuitive inference and epistemic vigilance stands out as a great candidate for solving the problem of dynamic choice.

Green, Alexander (University of Warwick)

Goal slippage: a mechanism for infant helping behaviour

Abstract: Over the past decade it has been well documented that children as young as eighteen months will spontaneously help others (Warneken & Tomasello, 2006; Warneken et al., 2007; Svetlova et al., 2010; Kenward & Gredeback, 2013; Barragan & Dweck, 2014; Hepach et al., 2017). The experimental setup in these studies is broadly similar: an experimenter is acting to bring about a particular goal, but then encounters an obstacle. Upon seeing this, children often spontaneously help.

Since spontaneous helping behaviour emerges so early in ontogeny, it may provide an important window into the development of moral cognition and prosocial behaviour. It is therefore no surprise that much of the research in this area has been devoted to identifying the mechanisms underpinning this behaviour.

The current paper investigates one possible mechanism which has recently been proposed to account for spontaneous helping in young children - goal slippage (Michael, Sebanz & Knoblich, 2016; Michael & Székely, 2017). Suppose that an agent is acting to bring about some goal and requires help. Goal slippage hypothesizes that children first identify the other agent's goal, and consequently form that goal themselves. That is, children take on the other agent's goal, such that this goal is now an outcome that the child wants to see brought about. What then motivates children to engage in helping behaviour is simply a desire to see their own goal satisfied.

While the goal slippage hypothesis is consistent with the available findings, it has yet to be directly tested. One reason for this is that it is difficult to distinguish goal slippage experimentally from the hypothesis that helping behaviour in young children is motivated by psychological altruism - i.e. by a concern to bring about the other agent's goal because it is the other agent's goal (Warneken & Tomasello, 2006; Warneken et al., 2007; Hepach et al., 2012; Barragan & Dweck, 2014; Hepach et al., 2017).

My aim is to disentangle these two hypotheses by spelling out an experimental design in which they lead to conflicting predictions about children's behaviour. The core manipulation involves the experimenter abandoning their goal. Under the psychological altruism hypothesis, children should be interested in the experimenter's goal because the experimenter is interested in the goal. If the experimenter abandons their goal, children should therefore lose their motivation for helping. In contrast, under the goal slippage

hypothesis, children should be interested in the experimenter's goal because they have taken on the experimenter's goal for themselves. Thus, the experimenter abandoning their goal should have little effect on infant helping rates.

In the presentation, I spell out an experimental design that can adjudicate between these two hypotheses. If results favour the goal slippage hypothesis, one important implication is that prosocial behaviour may emerge from more basic, selfish origins. If not, an important alternative to psychological altruism will have been ruled out. Either way, the results from an experiment implementing this design would shed important light on the mechanisms underpinning some of the earliest prosocial behaviour that we observe in humans.

Grodniewicz J.P. (University of Barcelona)

Linguistic understanding across communicative scenarios

Abstract: The paper confronts the current orthodoxy in the philosophy of language according to which we should ascribe linguistic understanding of a given utterance to a subject if and only if she is in a mental (or complex—mental plus environmental) state of some specific kind: knowledge, perception-like grasp, or content-entertaining. I argue that the common assumption of all these accounts is wrong. We will not be able to formulate a satisfactory characterization of linguistic understanding unless we observe that, in different communicative scenarios, different states may be equally aptly characterized as states of understanding.

My discussion proceeds in two steps. First, I summarize arguments presented in AUTHOR (XXXX) supporting the thesis that what is common to all instances of linguistic understanding is not a kind of hearer's cognitive state, but a specific cognitive process involved.

Second, analyzing arguments present in the debate, I argue that states resulting from this process vary across communicative scenarios. The variance takes place along at least two dimensions. First, examples provided by Hunter (1998), Pettit (2002) show (even though it was not the intention of the authors), that some states of understanding are belief states (*), while some others are mere phenomenal seemings (+). Second, arguments in (Longworth, 2016) show (again, contra Longworth), that some states of understanding represent in their content the speaker of the utterance (A), while some others do not (B). In result we obtain the following initial typology of states of understanding (where U stands for an utterance and S stands for a speaker):

(*A) belief that what was said in U by S was that p
(*B) belief that what was said in U was that p

(+A) seeming that what was said in U by S was that p
(+B) seeming that what was said in U was that p

All four are kinds of states of linguistic understanding as long as they originate from an appropriate cognitive process. What introduces the variance is the communicative situation the hearer is in (or her beliefs about the communicative situation). The state of understanding is of type (*A) or (*B) if the hearer does not have reasons to doubt her ability to understand given utterance. If she has reasons to doubt it (as in Pettit's and Hunter's mad scientist scenarios), she ends up in a state of the type (+A) or (+B). If the hearer

believes that she is in a highly cooperative or highly competitive scenario (in both of them it is important to track the identity of the speaker), she ends up in a state of the type (*A) or (+A). If in a given scenario the identity of the speaker does not matter or if the identity is unknown to the hearer, her understanding state belongs to the type (*B) or (+B).

References:
Hunter, D. (1998). Understanding and belief. *Philosophy and Phenomenological Research,* 58(3):559–580.
Longworth, G. (2016). Understanding what was said. *Synthese*, pages 1–20.
Pettit, D. (2002). Why knowledge is unnecessary for understanding language. *Mind*, 111(443):519–550.

Grünbaum, Thor (University of Copenhagen)

The Feeling of Agency: Paradigms and Model Testing

Abstract: Despite the plausibility of the claim that a feeling of agency is associated with particular movements, it has turned out be difficult to conceptualise and operationalise the target psychological construct (i.e. the phenomenal feeling of agency associated with particular voluntary movements). In this paper, I discuss a number of obstacles confronting contemporary experimental studies of the feeling of agency for movements. The main problem is that an apparent misalignment between the target psychological construct and standard experimental paradigms makes it difficult experimentally to test processing models (for instance, the comparator model) of the feeling of agency for movements.

Hare, Ian (University of East Anglia)

Low Conceptual Coherence: A New Hypothesis Concerning Autism

Abstract: This paper introduces a new hypothesis concerning autism: the low conceptual coherence (LCC) hypothesis. On this hypothesis, autistic people have less well integrated concepts, meaning they make weaker connections between feature representations in semantic memory. This results in autism symptoms via two mechanisms, concept weakening and concept hardening. Concept weakening means autistic people will make fewer automatic inferences than neurotypicals in the same situations. Concept hardening means they will take a relatively narrow number of cues into account in order to identify category members. I argue that weakening and hardening can jointly explain several of the traits that constitute autism, including social difficulties, unusual sensory experiences, and repetitive behaviours, as well as relative strengths in rule-governed domains.

Haukioja, Jussi (Norwegian University of Science and Technology), Mons Nyquist (Norwegian University of Science and Technology) and Jussi Jylkkä (Åbo Akademi University)

Reports from Twin Earth

Abstract: In this paper, we present results from two experiments that were conducted on ordinary speakers' usage of natural kind terms, using cases closely modeled after Putnam's classic Twin Earth case. The Twin Earth case, and other similar thought experiments, have widely been taken to provide strong support for an externalist view that combines a causal-historical theory of reference with essentialism: sharing of underlying nature with familiar samples of a kind is taken to be both necessary and sufficient for

belonging in the extension of a natural kind term. Previous experimental studies on natural kind terms have cast some doubt on this view, suggesting that ordinary speakers' usage of them displays a split pattern, where the usage is sometimes in accordance with the causal-historical theory, at other times with the descriptivist theory. However, the results of these studies are far from conclusive, and none of the existing studies have looked directly at Twin Earth type cases.

In our study, we presented subjects with two types of scenarios, using five different natural kinds (water, tiger, lightning, diamond, gold). The first type of scenario was a straightforward Twin Earth case, where new samples were found to share appearance, but not underlying structure, with the familiar samples of a kind (for example, a case where watery XYZ is found on a planet). The second type was a reverse case, where new samples were found to share underlying structure with familiar samples of a kind, but differ dramatically in appearance (for example, a case where solid and smelly H2O is found on a planet). The subjects' usage was probed both with an elicited production task, where they had to report on the new findings, and with forced-choice questions (where the relevant samples had to be classified as water or non-water, and similarly for other natural kinds). In addition, we attempted to see to what extent the subjects were inclined to defer to the relevant experts.

Our results were remarkably clear, and partly highly surprising. The standard Putnamean judgments concerning Twin Earth (for "water", as well as for other natural kind terms) were confirmed, but at the same time the subjects were systematically unwilling to apply a natural kind term to samples which shared an underlying structure, but not their appearance, with the familiar samples in the relevant kind term's extension. The subjects appear, then, to take both sharing of underlying structure and similarity of appearance with standard samples of a kind to be necessary for belonging to the extension of the relevant natural kind term. This result is quite problematic both for mainstream externalist and mainstream internalist theories of natural kind terms. The strong support for the standard Twin Earth judgements is a problem for traditional descriptivist views. However, contrary to externalist orthodoxy, sharing of underlying structure is treated as necessary, but not sufficient, for belonging in the extension of a natural kind term.

Hvorecky, Juraj (Czech Academy of Sciences)

Componential analysis of phenomenal consciousness

Despite the overwhelming support for the identity of phenomenality with consciousness, there have been few attempts to defend the existence of phenomenal properties at the unconscious level. Rosenthal (1999) has been the strongest defender of such a view. On his picture, unconscious mental states possess exactly the same properties that they manifest in consciousness. They remain unconscious until they are subject of a higher-order thoughts (HOT). Rosenthal's position has been recently supported by such diverse authors as Coleman (2015) and Marvan & Polak (2017). Coleman advocates a picture of the ascent of the full-fledged unconscious (proto-)phenomenality into consciousness via the process of "awareness". Marvan & Polak, employ various perceptual studies to demonstrate the existence of phenomenal states at the unconscious level.

We explore the sustainability of unconscious phenomenality. We show, pace Rosenthal, that unconscious phenomenality cannot possess the same qualities that we experience consciously. All the examples, presented by the defenders of unconscious phenomenality, are perceptual and come from a single modality. Yet consciousness is hardly ever single modal.

Examples of intermodal effects among conscious states (McGurk effect, spatial ventriloquism or rubber hand illusion) show that at the conscious levels novel phenomenality emerges that could not have existed unconsciously.

We then offer an alternative picture of unconscious phenomenality. We argue that instead of pointing out a single dimension of unconscious phenomenality, its defenders, in fact, point to a set of qualities that make up phenomenal content. Components of phenomenality include perspectivalness, mineness, intentionality, retention, intensity, unity and possibly other aspects. None of them is necessary or sufficient to bring about phenomenality on its own, but a various combination of these components produce phenomenal consciousness of content. Our componential analysis explains several crucial features of consciousness. First, it is highly likely that some components (i.e. perspectivalness, intentionality) are fully present at the unconscious level. Because of this, a mistake of assigning complete phenomenality to unconscious states might occur. Other components are more closely tied exclusively to conscious level (mineness, multimodal unity, intensity) and it is their combination with those already present at the unconscious level that makes up full-blown conscious phenomenality.Second, defenders of unconscious phenomenality envision a clear demarcation line between conscious and unconscious states. On the introduced componential picture, no strict dividing line exists. Whether a given state is conscious or not, depends on a presence of a bundle of components that might be dynamically changing and often are not very stable.

Finally, in-between conscious and unconscious states lays a grey zone where states can be marginally conscious, or conscious only upon executing a particular effort. We conclude by arguing that debates about the very notion of phenomenal properties are not particularly helpful because the notion is further analyzable.

Coleman, S. (2015). Quotational higher-order thought theory, *Philosophical Studies*, 172: 2705-2733.

Rosenthal, D. (1999). Sensory Quality and the Relocation Story. *Philosophical Topics*, 26: 321– 350.

Marvan, T. and Polak, M. (2017). Unitary and Dual Models of Phenomenal Consciousness. *Consciousness and Cognition*. 56, 1-12.

Hewson, Claire (The Open University) and Gabriela Saldanha (University of Birmingham)

An exploration of corpus analysis methods for investigating the explanatory framework of folk psychology

Abstract: Our everyday commonsense framework for explaining and predicting human behaviour – referred to as folk psychology – has been of interest to theorists working across diverse research areas and disciplines. The question of the key concepts and distinctions that underlie commonsense behavioural explanation has concerned psychologists, philosophers and cognitive scientists. However, cross-disciplinary perspectives on this question have remained largely disconnected, and have embodied contrasting emphases and claims. There remains a need to better integrate these respective bodies of research, and further clarify the nature of our everyday folk psychological explanatory framework. The present research contributes to this goal by demonstrating how corpus analysis approaches can help elucidate the central concepts and distinctions appealed to in folk psychological explanation. We present a study in which we selected and analysed a random sample of explanations found in naturally occurring, spoken and written, language and classified these in order to examine the type and prevalence of different explanatory

concepts and distinctions appealed to in folk psychology. In order to be able to consider and integrate existing competing perspectives, we employed a working definition of folk psychology that was suitably broad, namely: our everyday commonsense framework for describing, explaining and interpreting the behaviours and psychological states of ourselves and others. This allowed us to consider various perspectives that have been offered, including the prevalent view within philosophy and cognitive science of folk psychology as 'belief-desire talk', and the suggestion within social psychology that people fundamentally strive to attribute human behaviour to either 'person' or 'situation' causes, and to reduce variable behaviour to underlying stable dispositional properties (e.g. character, ability). Previous laboratory-based investigations have often assumed the distinctions being explored, thus pre-supposing these distinctions (e.g. person-situation) in their designs. Our study shows that explanations found in naturally occurring language can allow for more nuanced and, arguably, more accurate insights into the nature of folk psychology in ways not achievable using laboratory-based studies. We discuss how our findings derived from analysing natural language data compare with results from laboratory studies tackling similar questions. We also compare our findings with previous work in social cognition that has considered the types and prevalence of events that people choose to explain (i.e. the explananda of folk psychology). Then, we present our analysis of the explanans of folk psychology, considering the extent to which our findings support the various accounts that have been offered within the existing literature. We conclude by arguing for the integration of corpus analysis methods as as a valuable approach to exploring aspects of folk psychology, particularly the concepts and distinctions that underlie everyday commonsense behavioural explanation. We suggest possible future research avenues using this approach, and consider how it might usefully supplement existing empirical investigations of folk psychology. Gaining a clearer, empirically grounded picture of the conceptual framework underlying folk psychological explanation is important, since many further questions about folk psychology, such as its theoretic status and scientific credibility, rely on first gaining an accurate picture of quite what folk psychology is.

Ikier, Simay (Bahçeşehir University) and Nazim Gokel (Kilis 7 Aralık University)

Context Dependency and Memory

Abstract: Tulving's Encoding Specificity Principle (Tulving & Thomson, 1973) states that when information is encoded, it is encoded along with its context, and if this context is recreated at retrieval, memory performance is better. The Encoding Specificity Principle continues to generate empirical research and theoretical propositions, by using context in a multiply-defined way (Smith & Vela, 2001). The general experimental procedure in these studies involves a study session in which information such as a word list is encoded in one context (e.g., Context 1), and a test session in which the memory is retrieved either in the same (Context 1) or in a different context (Context 2), resulting in two matched (Context 1-Context 1; Context 2-Context 2) and two mismatched (Context 1-Context 2; Context 2-Context 1) conditions for encoding and retrieval. While the initial work on context-dependent memory concentrated on the context that is provided by the experimental conditions, such as the cues provided related to the studied items or the way that these items are processed at study and retrieval (e.g., at a superficial level vs. at a semantic level), later on, the idea of context was extended to the match between encoding and retrieval of the physical environment (e.g., whether the participants who are divers are under water vs. on land), the physiological state (e.g., whether the participants had alcohol vs. are sober), the psychological state (e.g. whether the participants are happy vs. sad), and linguistic context (e.g., whether the memory is being retrieved in the language that it was encoded in vs. another language). Smith and Vela (2001) provide four hypotheses about the strength of these context dependency effects: Reinstatement refers to increased context dependency when the incidental memory cues are reinstated at retrieval, while Overshadowing

and Outshining, refer to reduced context dependency when the processing of non-contextual information is encouraged at encoding or retrieval respectively, and Mental Reinstatement refers to the mental reinstatement of context, again resulting in reduced context-dependency. Their analyses support all hypotheses. Marian and Neisser (2000) conducted a study with Russian-English bilinguals who had immigrated to United States at an early age. The study found that the linguistic context in which memories for life events are encoded in increases the probability that these memories will be retrieved when the linguistic context is recreated at retrieval. Specifically, participants were more likely to retrieve memories from the periods they spoke Russian versus from the periods that they spoke English, when interviewed in Russian versus English respectively. The present paper aims to analyze linguistic dependency in memory within the framework that Smith and Vela (2001) provides and to identify when memory may be more language-dependent. Finally, the paper poses the question of whether memory processes are similar to thinking processes, and whether the Language of Thought hypothesis (Fodor, 1975) can be related to memory.

Irvine, Liz (Cardiff University)

Meaning and communicative success in interaction

Abstract: In this talk I use recent research on conversational interaction to identify a new empirical challenge to minimalism, and use this as a basis to challenge the explanatory targets of semantic theories more generally.

Minimalists claim that there are propositional, literal meanings of natural language sentences that are exhausted by the lexical and syntactic elements of those sentences, which are (almost entirely) free from contextual effects. Further, according to minimalists, "…it is possession of a theory of meaning which ultimately trades in sentence-level contents that explains (at least in part) why subjects are in a position to recover speaker meaning at all" (Borg 2012, p. 64). Minimalism is claimed to be empirically sensitive in a standard way: it helps to explain how successful communication is possible.

However, in light of recent work on the timing of conversations, it is hard to see what role sentence-level contents might play in identifying speaker meaning. Most communication occurs in interactive face-to-face conversations which exhibit rapid turn-taking behaviour. The time between turns is brief (200ms), so recipients must start planning their response before the speaker's turn has ended. To do this, all models of turn-taking assume that recipients predict speaker meaning early on during a turn (Levinson & Torreira 2015; Magyair & de Ruiter 2012). There is also evidence that infants take this predictive approach (Casillas et al. 2016).

This means that if recipients do compute the literal sentence contents of an utterance, it is not based on what the speaker actually says, but on the basis of a predicted utterance, in a fully specified propositional format.

The problem is that there is little reason to think that recipients do this. First, in models of turn-taking the prediction of syntactic structures and specific lexical content is at least partly based on prior prediction of speaker meaning. Indeed, it is not clear how else (for example) the specific lexical content of a sentence could be predicted. Second, predictions about specific features of an utterance are primarily used to accurately predict the end of a speaker's turn and so when to launch a prepared response. In this case, there is no reason to think that sentence-level content plays an explanatory role in how recipients recover

speaker meaning.

Yet in claiming to be empirically sensitivity, theories of semantics are aiming at an ecologically nonstandard explanatory target: successful communication at the sentence-level. Research on turn-taking shows that successful communication is about adequately responding to what you predict a speaker means, rather than accurately processing what they actually utter. Further, successfully understanding single utterances is surprisingly unimportant in self-evaluations of communicative success, providing participants follow norms about how to engage in conversational repair (Galantucci & Roberts 2014). Instead then, successful communication at the level of interactive conversations is a more ecologically plausible target by which to assess theories of meaning, and associated theories of communication. I will briefly sketch what explanatory tools are likely to be relevant to this task.

Jary, Mark (University of Roehampton)

Assertion and linguistics

Abstract: Grice sought to equate what is asserted by a speaker with the conventional meaning of an utterance, after reference assignment and disambiguation, i.e. his notion of 'what is said'. A number of authors have pointed out problems with this position, and have argued that what a speaker asserts goes beyond Grice's notion of what is said. This has led to the posting of notions such as explicature (Sperber & Wilson, 1995), impliciture (Bach, 1994) and enriched what is said (Recanati, 2002, 2004), hybrids of encoded meaning and pragmatically derived content.
The observation that what a speaker commits herself to generally goes beyond what she said, in the Gricean sense, has led theorists to propose models of utterance interpretation that aim at deriving this hybrid of encoded and inferred meaning. Thus, Relevance Theory posits that the cognitive processes of utterance interpretation are sensitive to the distinction between explicature and implicature, and assigns a special role in the interpretation process to explicitly conveyed content. Similarly, Recanati posits two distinct process, one concerned with what is said, in his pragmatically enriched sense, the other concerned with what is implicated.

I will argue that the distinction these authors draw is misguided. Rather than a simple distinction between explicit and implicit communication, I argue, we should first draw a distinction between what is behaviourally implicated and what is linguistically implicated (Jary, 2013). This can be done by means of rational reconstruction: to count as linguistically implicated, it must be possible to derive an implication as a reasonable conclusion in an inference from what is said, in the Gricean sense. If, by contrast, the initial premise required to derive an implicature must be of the form 'The speaker said that…', then the implicature is behavioural, and not properly linguistic. Once this distinction has been made, I will argue, there is no reason to posit that the cognitive processes involved in utterance interpretation are sensitive to the distinction between what is asserted (i.e. the truth conditional content of the utterance) and what is linguistically implicated. I will argue for this by showing that the role played by explicature in Relevance Theory is redundant. The result will be that basic pragmatic processes are not directly involved in identifying asserted content. Rather, assertion is a meta-communicative notion brought into play when issues of responsibility arise (cf. Stainton, 2016, Borg, 2017).

Jefferson, Anneli (University of Birmingham)

The case against adaptationist accounts of psychopathy

Abstract: In this paper, I criticise adaptationist accounts of psychopathy. Because antisocial, manipulative and exploitative behaviour is such a crucial part of the psychopathic personality profile, the objection that we are medicalizing a moral problem is commonly raised. Immoral behaviour for personal gain and instrumental aggression in order to achieve one's ends are clearly morally problematic. However, it is less obvious that they are detrimental to the individual behaving this way, especially if the metric of success is survival and reproduction, rather than some more rich notion of human flourishing. Adaptationist accounts argue that psychopathy was in fact selected for because it was beneficial to individuals exhibiting it. I discuss a number of adaptationist proposals and show why these are not plausible. Consequently, we should reject the claim that psychopathy is not a disorder because it is an adaptation.

Jordan, Maiya (University of Connecticut)

Self-Awareness and Inner Objects

Abstract: I defend Husserl's notion of pre-reflective self-awareness (or pre-reflection) against two objections. Kriegel has criticized Husserlian pre-reflection because, he claims, Husserlian pre-reflection does not apprehend experiences as objects. It is, therefore, mysterious. Conversely, Zahavi has urged that we should reject Husserlian pre-reflection if it does apprehend experiences as objects. I argue that each of these objections should be rejected. Contra Kriegel, Husserlian pre-reflection does apprehend experiences as objects. However, thus understood, Husserl's position does not succumb to Zahavi's criticisms.

Kärki, Kaisa (University of Jyväskylä)

Procedural metacognition in intentional omissions

Abstract: In this paper I argue that a necessary condition of an intentional omission is procedural metacognition – a kind of recognition by the agent of the possibility of action that is left undone. I argue that mental or physical effort are too strong conditions for intentional omission whereas mere guidance control is too weak. Because some mental activity is needed for an omission to count as intentional, in the most minimal sense, this mental activity is procedural metacognition through which the agent perveiced the possibility of their own action in the horizon of future actions. This is because without this metacognitive component, the agent cannot intentionally try not to do some action, resist doing an action or decide or intend not to do some action.

Koi, Polaris (University of Turku)

ADHD outside the genes: distinguishing etiology, ontology, and disadvantage

Abstract: Debates around what Attention Deficit / Hyperactivity Disorder is, how it is caused, and what our clinical and social responses to it should be like, revolve around claims concerning its ontology. The heritability of ADHD often plays a large role in these claims.

The paper is a contribution to a broader debate concerning the etiology and ontology of mental disabilities. Two archetypal positions, often called the medical model and the social model of disability, have served as signposts of the outer reaches of this debate. Few scholars defend the social model simpliciter, which states that disability is fully produced and explained by unjust social structures, and no competent medical practitioner would endorse the 'medical model', according to which disabilities are physical impairments and any harm resulting from them is due to the medical condition of the disabled individual. While both

models may appear overly simplistic, both popular and, to an extent, academic discourses continue to represent ADHD in terms of a corresponding nature / nurture debate. While this debate is especially polarized in the case of ADHD, a similar conceptual framework surrounds mental disorders and disabilities in general.

It is true that studies into the behavioural and molecular genetics of ADHD have shown it to be highly familial and heritable. However, these findings are often simplistically interpreted as proving that ADHD is primarily a neurobiological rather than a social phenomenon in its ontology. These notions then shape how people with ADHD are encountered, and how they and their families negotiate for recognition. In other words, ideas about the etiology of ADHD shape positions on the ontology of ADHD, which in turn impact what we believe to be a proper response or a fitting form of care for people with ADHD diagnoses, and are employed to defend a variety of stances on the matter. Put briefly, the etiology, ontology, and social disadvantage are often bundled up as a single explanandum, to be solved with a single explanans.

In this paper, I disentangle etiology, ontology, and social disadvantage in ADHD. I suggest that in addition to philosophical methods, concepts in behavioural genetics can help us make important distinctions between these muddled categories, distinctions that are not only of ontological interest but also of ethical and clinical relevance.

Drawing from studies in behavioural genetics, I demonstrate that its ontology or the ontological features of the resulting disavantage cannot be deduced from its etiology, and describe some attempts to account for the interplay of the social and biological aspects of mental disabilities and the problems they run into when faced with explaining ADHD.

Kopecký, Robin (Charles University) and Michaela Košová (Charles University)

How virtue signalling makes us better: Moral preference of selection of types of autonomous vehicles

Abstract: In this paper, we present a study on moral judgement on autonomous vehicles (AV). We employ a hypothetical choice of three types of "moral" software in a collision situation ("selfish", "altruistic", and "aversive to harm") in order to investigate moral judgement beyond this social dilemma in the Czech population we aim to answer two research questions: Whether the public circumstances (i.e. if the software choice is visible at the first glance) make the personal choice "altruistic" and what type of situation is most problematic for the "altruistic" choice (namely if it is the public one, the personal one, or the one for a person's offspring).
We devised a web-based study running between May and December of 2017 and gathered 2769 respondents (1799 women, 970 men; age IQR: 25-32). This study was a part of research preregistered at OSF before start of data gathering.
The AV-focused block of the questionnaire was opened by a brief information on AV and three proposed program solutions for previously introduced "trolley problem like" collisions: "selfish" (with preference for passengers in the car), "altruistic" (with preference for the highest number of saved lives), and "aversion to harm" (which will not actively change direction leading to killing a pedestrian or a passenger, even though it would save more lives in total). Participants were asked the following four questions: 1. What type of software would you choose for your own car if nobody was able to find out about your choice ("secret/self"). 2. What type of software would you choose for your own car if your choice was visible at the first glance ("public/self"). 3. What type of software would you choose for the car of your beloved child if nobody was able to find out ("child"). 4. What type of software would you vote for in secret in the

parliament if it was to become the only legal type of AV ("parliament").

The results are as follows, test of independence was performed by a chi square: "Secret/self": "selfish" (45.2 %), "altruistic" (45.2 %), "aversion to harm" (9.6 %). "public/self: "selfish" (30 %), "altruistic" (58.1 %), "aversion to harm" (11.8 %). In public choice, people were less likely to choose selfish software for their own car.

"Child": "selfish" (66.6 %), "altruistic" (27.9 %), "aversion to harm" (5.6 %). A vote in parliament for legalizing single type: "selfish" (20.6 %), "altruistic" (66.9 %), "aversion to harm" (12.5 %) In choice of car for one's own child people were more likely to choose selfish software than in the choice for themselves. Based on the results, we can conclude that the public choice is more likely to pressure consumers to accept the altruistic solution making it a reasonable and relatively cheap way to shift them towards higher morality. In less favourable news, the general public tends to heightened sensibility and selfishness in case of one's own offspring, and a careful approach is needed to prevent moral panic.

Košová, Michaela (Charles University) and Robin Kopecký (Charles University)

The True Self is Good: Children's Concept of Personal Identity

Abstract: This paper addresses recent x-phi research on children's concept of personal identity and its relation to moral traits and interpersonal relationships. It is inspired by earlier studies showing that there is a strong connection between the folk concept of personal identity and preference for moral traits with interpersonal connotations (the "essential moral self" hypothesis of Strohminger & Nichols, 2014, 2015). Strohminger, Knobe and Newman (2017) suggest that the "true self" is generally viewed to be moral and inherently good. Also Tobia (2015, 2016) found that most respondents have higher tendency to agree that personal identity was broken after negative moral change than after positive moral change. These and other findings suggest that there is a connection between personal identity and positive interpersonal dispositions.

In order to explore these concepts also in developmental context we conducted an interview study in 2017 on Czech children and teenagers (N=217; 56,4% female; age range 6-15; average age=11). Respondents were randomly recruited at a public family event. Interviewer introduced each participant to a scenario in which a person ("your friend" (N=90), "someone you know" (N=36) or "some person" (N=91)) undergoes various changes after being closed in a special sci-fi chamber. Changes encompassed 6 categories: physical (appearance), cognitive (intelligence), moral (love for others, treatment of others), in character (laziness), in memory (remembering life experiences) and in perception (vision). Both negative and positive versions of the changes were included. Respondents were asked to judge how much each of the changes would affect the person's identity core on a 7-point scale (0-they are still the same person; 6-they are not the same person anymore).

Data analyses showed that respondents consider moral traits to be significantly more important for personal identity preservation than any other category of traits. The next most important category is memory, followed by cognition, character and perception. Physical category remains far behind all the other categories.

The expected difference between the impact of negative and positive versions of the changes was also supported by the data. The most salient difference came up in rating of the change in treatment of other people (becomes cruel vs. becomes nicer to other people). Also memory loss was rated considerably higher than the corresponding super-memory gain. On the other hand, negative and positive versions of the

change in physical appearance (becomes uglier vs. becomes more beautiful) got almost equal rating.

Further exploratory analyses revealed effects of age and scenario. Relative importance of moral traits grows with age, especially the distance between moral and physical category. Children responding to the personal scenario ("your friend/someone you know changes") ascribe higher importance to moral traits in comparison to other categories, than children responding to the neutral scenario.

We can conclude that the "essential moral self" and the "true self" concepts seem to be present already in children and become more salient with age and in certain personal contexts (friendship). These findings again point to a strong conceptual connection between personal identity and its impact on the quality of interpersonal relationships.

Kraft, Tim (University of Regensburg)

Three experimental approaches to epistemic closure

Abstract: According to the epistemic closure principle, if someone knows that P and also knows that P entails Q, she must know that Q as well. Although this principle is central to many epistemological debates and is often defended by relying on its intuitiveness, it has been put to the empirical test only recently. After Turri (2015) found that closure is violated in folk knowledge ascription surprisingly often, Beebe & Monaghan (2018) and [blinded] (2018) came to the opposite conclusion: By and large folk knowledge ascriptions adhere to the closure principle. Thus, there are now three experimental philosophy papers discussing whether closure is a principle of folk epistemology, but with different methodologies and different results. This is a fortunate and fascinating situation because it allows us to deepen our understanding of folk epistemology by comparing different studies. The aim of this paper is to present these three studies, to compare their methodologies and to look at what overall picture is drawn by them. My tentative result will be that the research conducted so far suggests that closure is a principle of folk epistemology after all.

All three papers rely on variations of Vogel's famous car theft case: Can you know that your car is where you parked it an hour ago if you don't know that it hasn't been stolen since then? Particpiants who respond that someone can know where one's car is without knowing that it hasn't been stolen violate the closure principle. Turri's original study raised two worries: Responses may be driven by pragmatic effects ("has been parked in" vs. "is now at") and the role of the manipulation used (epistemic source: seeing vs. inferring) in the argument against folk closure is unclear. Beebe & Monaghan and [blinded] improve on Turri's original study in interesting, but diverging ways. While Beebe & Monaghan rely on modified vignettes, a different question format (Likert-scales) and lay and expert samples (mathematicians who shouldn't overlook the relevant entailments), [blinded] rely on a completely different kind of manipulation (entailment instead of epistemic source). In this paper I discuss four questions I take to be central for interpreting these experimental results on folk closure:

1. What are the differences between the vignettes used and what is the rationale behind these differences?

2. What are the advantages and disadvantages of different question formats (Likert-scale vs. forced choice) in this context?

3. What manipulation is used and why?

4. How do the results fit with the non-experimental literature on Vogel-style cases (Douven 2007, Roush 2010, Nagel 2011) and the psychological claims made by them?

Finally, I will discuss some open questions and directions for further research.

Kraš, Tihana (University of Rijeka) and Maja Miličević Petrović (University of Belgrade)

Interpretation of Italian subject pronouns in Croatian-Italian simultaneous bilinguals

Abstract: Properties at the discourse-syntax interface have proved problematic in bilingual L1 and L2 acquisition and L1 attrition. This has been shown in particular for the use and interpretation of subject pronouns in null subject languages. Compared to monolinguals, highly proficient bilinguals, especially speakers of a null- and a non-null-subject language, tend to over-accept and overuse infelicitous overt pronouns referring to topical antecedents, while being (mostly) target-like on null pronouns. Two broad explanations have been proposed. The representational account (Tsimpli, Heycock, & Filiaci, 2004) attributes the bilinguals' problems to cross-linguistic influence, while the processing account ascribes the difficulties primarily to bilinguals' hypothesised less-than-optimal processing abilities (Sorace & Filiaci 2006). According to the representational account, difficulties should not arise when two grammatical systems pattern together with respect to an interface property; according to the processing account, difficulties should occur even then.

To test the predictions of the two accounts, we conducted an experimental study into the interpretation of Italian subject pronouns in intra-sentential contexts by Croatian-Italian simultaneous bilinguals (N=40) and a control group of Italian monolinguals (N=48), aged 11-15. The two languages involved pattern together with regard to the antecedent biases of null and overt subject pronouns. Participants read sentences containing null and overt pronouns, which either followed (Il testimone indica l'accusato mentre pro/lui entra in tribunale 'The witness points to the accused as he enters the courtroom') or preceded the candidate antecedents (Mentre pro/lui entra in tribunale, il testimone indica l'accusato 'As he enters the courtroom, the witness points to the accused') (anaphora vs. cataphora), and matched each sentence to one of three pictures, showing the antecedent as the matrix subject, the matrix object or an extra-linguistic referent. The task is an adaptation of task used by Tsimpli et al. (2004) and Sorace and Filiaci (2006), the results of which, based on off-line measures, have provided the basis for the two accounts.

The bilinguals expressed the same antecedent preferences as the monolinguals in all conditions apart from cataphora with overt pronouns, where they chose the topical, subject antecedent less often than the monolinguals, i.e. in 24.1% compared to 37.5% of the cases. In other words, it was the monolinguals, rather than the bilinguals, who accepted more overt pronouns referring to discourse topics. However, the difference did not reach statistical significance in a logistic regression analysis. We compare the results with previous findings in the L2 acquisition domain and interpret them as pointing to cross-linguistic influence, and thus lending support to the representational account.

References

Sorace, A., & Filiaci, F. (2006). Anaphora resolution in near-native speakers of Italian. Second Language Research, 22(3), 339-368.

Tsimpli, I.-M., Sorace, A., Heycock, C., & Filiaci, F. (2004). First language attrition and syntactic subjects: A study of Greek and Italian near-native speakers of English. International Journal of Bilingualism, 8(3), 257-277.

Krzyzanowska, Karolina (University of Amsterdam), Peter Collins (Ludwig Maximilian University of Munich) and Ulrike Hahn (University of London, Birkbeck)

Learning Probabilistic Relevance from Indicative Conditionals

Abstract: Our beliefs change in response to what we learn, and a lot of what we learn comes from the testimony of other people. The extent to which our beliefs change may depend on how many people are the sources of given information, or how reliable their expertise makes them be. It is not entirely clear, however, what the effects of reliability or the number of speakers will be when the testimony has the form of an indicative conditional. In this paper, we test a hypothesis that learning that if p, then q amounts to learning that p is a reason for q, or, more formally, that p raises the probability of q. We investigate whether this effect depends on the number and the expertise of speakers asserting the conditional.

Laasik, Kristjan (Zhejiang University)

Phenomenology and Perceptual Content

Abstract: In their seminal paper "The Intentionality of Phenomenology and the Phenomenology of Intentionality", Terence Horgan and John Tienson argue for the following claim,
Phenomenal Intentionality: There is a kind of intentionality, pervasive in human mental life, that is constitutively determined by phenomenology alone" (Horgan and Tienson 2002, p. 520).
The phrase "constitutively determined" means that intentional states have such content by virtue of their phenomenology. More specifically, they purport to show that phenomenal intentionality is a feature of our perceptual experiences and some propositional attitudes. Since my focus in this paper will be on perceptual experiences, I will generally not distinguish between Phenomenal Intentionality (PhenInt) and the idea that our perceptual experiences have phenomenal intentionality.
Horgan and Tienson's arguments crucially involve a phenomenal duplication thought experiment, inviting us to accept that, should there obtain a suitable pair of phenomenal duplicates, i.e., subjects whose experiences have exactly similar phenomenology, then these experiences would have the same intentional content. However, Bailey and Richards object that, in order for PhenInt to follow, Horgan and Tienson would first have to establish the co-variation of phenomenology and intentional content, rather than just the uni-directional determination of intentional content by phenomenology (Bailey and Richards 2014). Moreover, even if they succeeded in establishing co-variation, PhenInt would still emerge as less plausible than its converse, representationalism, i.e., the view that intentionality is more basic than phenomenology, and that phenomenology depends on, or is reducible to, intentionality.
These objections, pitting PhenInt against the more mainstream representationalist alternative, are certainly a serious challenge that needs to be taken up by defenders of PhenInt. I will therefore address the challenge by appeal to Husserlian ideas, centered on the notions of fulfillment and constitution. Doing so will involve suitably clarifying PhenInt, as well as modifying Horgan and Tienson's argument. I will propose that we regard phenomenology, or what it is like to undergo a certain experience, as comprised of proto-phenomenology and phenomenology proper. By proto-phenomenology I mean non-intentional phenomenology, and I centrally have in mind visual and other sensations. When I speak of phenomenology proper, I mean the complex, intentional phenomenology of everyday experience. I will, on the one hand, argue that phenomenology proper determines a kind of content that is conceived in terms of fulfillment conditions, or what it takes to bring aspects of objects and scenes to different, and better, givenness, in the further course of the intentional experience. I will, on the other hand, argue that we can establish the primacy of phenomenology, relative to intentional content, by tracing back the phenomenology proper and

the intentional content to the pertinent proto-phenomenology, which functions to "constitute" it—the Husserlian term for the achievement of a kind of experiential unity, based on the functioning of visual and other sensations. However, my main aim is not to invoke Edmund Husserl's views as they stand, but rather to explore which aspects of the Husserlian view can help us make progress in the phenomenal intentionality debate.

Löhr, Guido Robin (Ruhr Universität Bochum and Institut Jean Nicod)

Polysemy, co-predication and minimalism: why we need a new kind of moderate Contextualism

Abstract: Minimalists and most current moderate contextualist positions, which I call "modulation theories", make one crucial assumption that I call "the literalist principle": linguistic processing takes lexical concepts (the concept associated with the expression to be comprehended) as input for further primary or secondary pragmatic processing. If the literalist principle is false, then minimalism and most modulation theories are false. I argue that the literalist principle gets linguistic processing the wrong way round. Early linguistic processing involves the activation not of concepts or meanings, but of context-dependent collections of information that include prototypes and exemplars. These bundles of information are indeterminate as to what concept or meaning they are used to express. Speakers often commit only in a more complex post-hoc process to a content or concept. The rejection of the literalist principle leads to a new theory of polysemy for psycholinguistics that I call "the words as evidence approach to polysemy". This view renders minimalism and modulation theories redundant and illustrates why they are incompatible with the empirical evidence from psychology.

Lorge, Isabelle (University of Cambridge) and Napoleon Katsos (University of Cambridge)

Theory of mind or cognitive biases? Factors influencing pragmatic inference and audience design

Abstract: By definition, pragmatic inferences or implicatures which are arrived at by assuming the speaker to be informative are not always derived. Similarly, speakers do not always design their utterances to be optimally informative or in a hearer-oriented way (Keysar et al., 2000). Other cognitive factors and biases such as salience or default/automatic responses and procedures come into play. Some responses and behaviours resulting from these biases have been given higher-level interpretations which might actually not reflect psychological reality or indeed the mechanisms typically called for in similar situations.

We first present a word learning task inspired from Frank and Goodman's (2014) where participants are shown two characters (e.g., kittens) with novel objects. One of the characters has object A whereas the other has both object A and B. Below are two frogs, one with object A only and one with object B only. The instruction states 'Oh! A kitten with a fep! Can you touch the frog that has a fep?'. The expected implicature is to assume that the speaker intends to be informative and would not draw attention to an object shared by both characters, thus choosing the kitten with the unique object (B) and the corresponding frog.
Developmental effects: In a version where adults and children had to guess which kitten was being targeted, adults are at chance (m=0.5, t=0, df=127, p=1) whereas children seem to be more affected by the similarity between the image of the kitten with only object A and the frog with the same object and choose frog with object A significantly more than chance (m=0.65, t=-6.29, df=416, p>0.0001).
Salience dilution effect: We ran two versions of the experiment where adults were explicitly told which

kitten was being targeted. In the first version adults were deriving the implicature m=0.64, t = 3.1411, df = 111, p<0.01)., but when a familiar object was added to the second kitten for reasons of symmetry, which should not have made a difference, adults were at chance again (m=0.48, t = -0.42441, df = 87, p= ns).

We ran another experiment where participants had to describe to an imaginary interlocutor an object to be picked between two options (e.g., a small and big blue banana)
Default audience design: In production, typicality effects have been attributed to audience design (Kreiss et al., 2017) ; i.e., mentioning the colour in 'blue banana' is helpful even when there is only one since it should help the interlocutor identify the unusually coloured object. However, in our experiment, even when the modifier was explicitly not helpful (two blue bananas or green bones), participants still produced them significantly more when the modifier was atypical than when orthogonal (e.g., blue car) (t = 2.671, df = 103, p = 0.008) or typical (e.g., yellow banana) (t= 4.4336, df = 103, p<0.0001)

We discuss the result of these experiments in terms of default and salience effects as well as the methodological manipulation of cognitive effort, rewards and feedback influencing the activation of theory of mind processes.

Strasser, Anna (Humboldt-Universität zu Berlin)

Beware of dichotomous conceptions

Abstract: It seems to be an open question what belongs to the realm of cognition and how one can specify what is meant by less cognitive processes. Standard cognitivist approaches focus on demanding representational information processes of adult human beings. However, interdisciplinary debates include abilities of children and non-human animals for which we need less-demanding notions. So-called minimal approaches are promising candidates to deliver appropriate notions. Focusing on mindreading abilities, I want to explore whether a characterization along the ideas of a two-system approach provides an appropriate description of cognitively less-demanding processes. This investigation will show that an dichotomous understanding of the two-system approach fails to capture the diversity of cognitive processes.

Luvisotto, Giulia (University of Warwick)

Moran's Authority and Estrangement: Epistemic or non-epistemic explanation of first-person authority?

Abstract: The present paper aims to understand how Richard Moran's view can account for first-person authority, conceived as the presumption of veridicality usually bestowed to present-tense self-ascriptions of psychological attitudes. Traditionally, philosophers have given epistemic or non-epistemic explanations. The former approach explains the authority by resorting to the speaker's special kind of knowledge of her attitudes, whereas the latter avoids to appeal to the speaker's knowledge as bearing any explanatory power.

On the face of it, Moran seems to be putting forth a non-epistemic solution by appealing to the notion of 'practical self-knowledge'. The present paper aims to evaluate whether this is a viable option and argues that either we admit that an epistemic element is unavoidable or we need to reshape our understanding of the 'authority' we wanted to account for.

More in particular, Moran makes a negative and a positive point. He overtly rejects the traditional epistemic explanation, arguing that it fails to capture the distinctiveness of the first-person perspective: self-knowledge cannot be a matter of privileged epistemic access. Hence, he rather puts forth the notion of 'practical self-knowledge' as capable of explaining non-epistemically first-person authority. More precisely, one would be authoritative as the only one in the position to determine one's attitudes simply by reflecting on the reasons for or against them. However, problems emerge on both sides.

First, Moran assumes but does not account for three conditions of possibility that the achievement of such deliberative self-knowledge would require, namely:
a. that I know what I hold to be true;
b. that I know that what I hold to be true is a belief of mine;
c. that I conceive of myself as a rational agent.
Nonetheless, especially the first one seems to render the proposal an epistemic solution: deliberating would still give one a kind of access to what one holds to be true. even not in the traditional 'observational' sense.

Second, it does not seem to follow from that fact that I am the only one capable of determining my attitudes the fact that my self-ascriptions are more likely to be true. In fact, to self-ascribe an attitude, it is necessary for me also to know what I am doing, i.e. that I am determining it (for instance, to self-ascribe "I believe that p", I must know that I am committing myself to the truth of p). Therefore, either we admit that an epistemic element is indispensable to account for first-person authority, or we modify our understanding of 'authority', so that one is not authoritative in that her self-ascriptions are more likely to be true but in that one is supposed to take on responsibility for whatever attitude she happens to determine, regardless her being aware of it.

Macià, Josep (University of Barcelona)

Presupposition and Metalinguistic Discourse

Abstract: Providing a proper characterization of what presuppositions are should help elucidate the nature of meaning and of the different levels, aspects or kinds of meaning (asserted meaning, literal meaning, character, content, conversational implicatures, explicatures, presuppositions, conventional implicatures, expressive meaning, etc.). There is a promising view of presuppositions which treats them as semantic entailments, and which models them using partial meaning functions. This promising view is challenged by the views put forward by those authors (for instance Soames 1989, Green 2000, Potts 2005 and Abbott 2008) that argue that presuppositions (like conversational implicatures) are cancellable. I will argue that one main kind of data that has been used to argue for this claim can be appropriately accounted for assuming the semantic view of presuppositions. I will, therefore, lend support to the view that we have no good reason to abandon the simpler, clear, promising semantic view according to which presuppositions are appropriately modelled using partial meaning functions.
As part of my argument, I will present a new kind of data that provides the basis for what I will call the "Fumfling company reply" and I will introduce and discuss the usefulness of the notions of "metalinguistic discourse" and of "remake of contexts".

Mahr, Johannes (Central European University) and Gergely Csibra (Central European University)

The importance of testimony: the role of epistemic authority about past events in human communication and sociality

Abstract: Unique events in our personal past seem to have a special significance for human beings. What grounds the special status of such events? It is commonly assumed that the main way we relate to the past is through the capacity of episodic memory. Episodic memory, however, differs crucially from other memory systems. When we remember episodically, we do not simply 'know' what happened, we also 'know how we know' what happened, namely that we experienced it. In other words, episodic memory allows us to become witnesses of the past. The question we want to address here is why humans have developed such a metacognitive memory system for the representation of unique past events. In other words, what is it about past events that humans developed the capacity to function as witnesses and give testimony about them? We will argue that because unique past events ground a large range of social facts (such as commitments and entitlements), they require a dedicated capacity allowing us to negotiate them effectively in communication. For other species it is similarly true that reciprocal relations are based on unique events, which requires that agents keep score of who did what to whom. This however, does not require the capacity to remember unique events as such. Because humans depend to an extraordinary extent on the capacity to influence and be influenced by others through communicative interaction in belief formation, we require a dedicated capacity allowing us to handle events as reasons in making and scrutinizing claims about how commitments and entitlements should be distributed. While, in principle, the availability of unique past events in a metarepresentational format allows us to use them in the transmission of generic beliefs and token causal judgments, we will argue that – for humans – communication about specific events is particularly important in so far as it allows the coordination of social facts and the policing of the social commitments and entitlements we are engaged in.

Marchi, Francesco (Ruhr University Bochum) and Albert Newen (Ruhr University Bochum)

Self-deception in the predictive mind

Abstract: In a study of a million US pupils, 70% thought that their abilities were better than average. Data such as this suggest that self-deception is a ubiquitous phenomenon that we need to include in our model of human cognition. Self-deception consists in forming and retaining a belief despite counter evidence, as a result of a motivational bias. We examine whether or not one of the currently most popular cognitive models, namely predictive processing, can account for this everyday phenomenon. Predictive processing rests on the assumption that biological systems behave efficiently and rationally by minimizing prediction-error arising from the divergence between a system's current and expected states, according to Bayesian principles. However, in self-deception it seems that the cognitive system refuses to update its model in the light of the prediction error that arises from testing a false hypothesis in the light of counter evidence. The core question is thus the following: How can predictive processing account for self-deception? Our account demonstrates how this is possible and uncovers that the prediction-error minimization framework is currently lacking a convincing integration of indispensable motivational factors in reasoning and perceiving. We close by offering a partial solution to this additional demand.

Martínez, Manolo (University of Barcelona)

Representations are Rate-Distortion Sweet Spots

Abstract: Information is widely perceived as essential to the study of communication and representation; still, theorists working on these topics often take themselves not to be centrally concerned with "Shannon information", as it is often put, but with some other, sometimes called "semantic" or "nonnatural", kind of information. This perception is wrong. Shannon's theory of information is the only one we need.

I intend to make good on this last assertion by canvassing a fully (Shannon) informational answer to the metasemantic question of what makes something a representation, for a certain important family of cases. This answer and the accompanying theory, which represents a significant departure from the broadly Dretskean philosophical mainstream, will show how a number of threads in the literature on naturalistic metasemantics, aimed at describing the purportedly non-informational ingredients in representation, actually belong in the same coherent, purely information-theoretic picture.

Marvan, Tomas (Czech Academy of Sciences)

Unitary, dual and hybrid theories of consciousness

Phenomenal consciousness "is the property mental states, events, and processes have when, and only when, there is something it is like for their subject to undergo them, or be in them" (Uriah Kriegel, Consciousness, Theories of, Philosophy Compass 1/1, 2006, p. 58). This widely accepted definition of phenomenal consciousness can be read in different ways. An assumption shared by most philosophers of mind is that phenomenality and consciousness are inseparable. That is, according to this received view, the qualities constituting the phenomenal character of a mental state can only occur at the level of consciousness. Let us call this theory of consciousness "unitary", because it takes the production of a phenomenal quality and it's becoming conscious for a subject to be one and the same thing. However, the assumption that phenomenality and consciousness are inextricable can be questioned. The "dual" theory of consciousness distinguishes the process in which the phenomenal quality is formed from the process that makes this quality conscious for a subject. It permits a separation of phenomenality from consciousness and, moreover, allows for non-conscious phenomenality. The dual theory was defended by David Rosenthal and some of his collaborators (such as Benjamin Young and Andreas Keller); cognate ideas can be found in the writings of Sam Coleman and Dimitris Platchias. Finally, according to the "hybrid" theory of consciousness we can in principle separate consciousness from phenomenality, but in reality the two always co-occur (Jesse Prinz defends such a theory in his 2012 book The Conscious Brain). I call this theory hybrid because it combines aspects of both unitary and dual theory. It is similar to the unitary theory in that it rejects non-conscious phenomenality but it shares the insistence of the dual theory that notions of consciousness and phenomenality need to be kept separate in the theory of consciousness.

I favor the dual theory, but my aim in this presentation is not to argue for it in detail. I want to delineate the fundamental commitments of the three approaches to consciousness and to compare their relative strengths and weaknesses. In particular, I will be showing that although the unitary theory is in deep agreement with widespread intuitions about consciousness, it is not in deep agreement with empirical consciousness research. In particular, it is unclear how it copes with some well-explored phenomena such as unconscious perception and perceptual priming. The unitary theory is forced to postulate two different mechanisms for conscious and unconscious perception, and is thus quite unparsimonious. This weakness of

the unitary theory is shared by the hybrid theory, but the hybrid theory avoids some of the counter-intuitive aspects and consequences of the dual theory that make the dual picture hard to sell. On the other hand, the dual theory seems to be not just compatible, but actually supported by some of the most popular hypotheses in the contemporary cognitive neuroscience of consciousness.

--

Matuszkiewicz, Maria (University of Warsaw)

Does singularity name a hybrid or a non-hybrid psychological category?

Abstract: Singularity is a semantic notion used to characterize both linguistic entities and thoughts. In my paper I will focus on the question whether there is a psychological property that corresponds to this semantic notion. First, I will distinguish three different questions: (i) whether the predicate singular thought names a natural kind, (ii) whether it names a hybrid or a non-hybrid category and (iii) whether it names an ascriber-dependent or an ascriber-independent category. I will relate this distinction to another, and partly orthogonal one: (a) whether singularity is an internal property of thoughts, or (b) whether it is relational and ascriber-independent one, or (c) whether it is relational and asciber-dependent one.
After proposing a brief typology of the views concerning psychological or epistemic conditions on singular thoughts, I will focus on two recent proposals, which define singularity in terms of a syntactic feature (Jeshion 2010; Recanati 2010, 2012). According to these views one is thinking singularly iff one is thinking through a mental file. Jeshion and Recanati provide different accounts of what makes this syntactic feature a hallmark of singularity. Jeshion's account draws on recent studies in perception and on the notion of visual indexes and object files postulated, among others, by Pylyshyn (2007). To explain singularity of mental files in the case of non-perceptual thoughts Jeshion makes one controversial assumption. Recanati accounts for mental files in terms of two triggering mechanism (acquaintance and decoding of language). Jeshion's proposal attempts to save an internalist intuition that thinking singularly does not depend on matters which are external to the thinker: it does not require any causal relation between the thinker and the object, nor does it even require that there exists a referent of one's thought. At the same time thinking singularly cannot be exercised merely at will: the object needs to be significant to the thinker. Recanati on the other hand defends the acquaintance requirement, although he qualifies it in an important way making acquaintance a normative, rather than factual constraint of singular thoughts. Contrary to Jeshion he holds that one cannot think singularly when there is no referent of one's thoughts.
I will conclude reflecting on the merits and the weaknesses of both proposals and by linking it to the initial question: whether singularity names a hybrid or a non-hybrid psychological category. Finally, I will sketch a more deflationary view according to which singularity is a property of contents, but contents are merely ways of characterizing psychological states, rather than properties of those states.

Mazzarella, Diana (University of Neuchâtel) and Filippo Domaneschi (University of Genoa)

Presuppositional effects and ostensive-inferential communication

Abstract: Philosophers and linguists have long debated the phenomenon of 'presupposition', whereby speakers linguistically mark information as background or taken for granted. This debate has surprisingly left untouched the following two assumptions: (i) presuppositions are not part of what is intentionally communicated by the speaker, and (ii) the propositional content of a presupposition is semantically determined. These assumptions are well-established in the literature and they are endorsed even by scholars who offer a 'pragmatic' analysis of presuppositions. For instance, with regard to (i), Simons (2007) claims that presuppositions are "not part of the speaker's communicative intention" and they are transmitted simply as a "by-product". Or, with regard to (ii), she suggests that the presuppositional content is "calculated, presumably compositionally, on the basis of the content of the trigger plus the rest of the content expressed" (Simons, 2005).

In this paper, we challenge both these assumptions. We argue that presuppositions fall within the scope of 'ostensive-inferential communication' (Sperber & Wilson, 1986/1995) and illustrate the benefits of our proposal. Ostensive-inferential communication requires the communicator to display an overt behavior aimed at attracting attention to her communicative intention ('ostension'). Furthermore, it requires the audience to infer the communicator's intended meaning ('inference').

First, we show that by treating presuppositions as part of what is ostensively communicated by the speaker, we can provide a unified account of a variety of presuppositional uses discussed in the literature. This account covers the whole range of cases from noninformative, or 'common ground' presuppositions, to informative presuppositions, as well as cases of exploitative presuppositions.

Second, we suggest that by treating presuppositions as the output of an inferential process of pragmatic interpretation, we can shed a new light on their context-sensitivity. Traditional accounts of presuppositions view the context-sensitivity of (at least some) presuppositions as pertaining to the question of whether presuppositions are contextually defeasible. This question does not exhaust the relation between presuppositions and context. The propositional content of a presupposition is the result of a process involving both semantic decoding and pragmatic inference, and involving a 'mutual parallel adjustment' (Wilson & Sperber, 2004) among presuppositions, the explicit content and the implicatures of the utterance.

Ostension and inference represent two sides of the same coin and can shed a complementary light on presupposition as a genuinely communicative phenomenon.

REFERENCES
Simons, M. (2005). Presuppositions and relevance. In: Szabo, Z. (Ed.), Semantics Versus Pragmatics. (pp. 329-355). Oxford: Oxford University Press.
Simons, M. (2007). Presupposition and cooperation. Carnagy Mallon University, ms, https://www.cmu.edu/dietrich/philosophy/docs/simons/Presupposition%20and%20cooperation.non-anonymous.pdf
Sperber, D., & Wilson, D. (1986/1995). Relevance: Communication and Cognition. Oxford: Blackwell.
Wilson, D., & Sperber, D. (2004). Relevance Theory. In L. R. Horn & G. L. Ward (Eds.), Handbook of Pragmatics (pp. 607–632). Oxford: Blackwell Publishing.

Michael, John (University of Warwick)

The Chains of Habit: Repeated Coordination in Joint Decision-Making Elicits a Sense of Commitment

Abstract:

The prevalence and flexibility of human cooperation is unparalleled by any other species. We routinely work together to achieve ends that we could not achieve alone, even setting aside short-term interests to maximize the benefits to our interaction partners and larger social groups.

In recent decades, a great deal of research in evolutionary theory, experimental economics and psychology has been devoted to investigating the evolutionary origins of human cooperation. This has led to significant progress in specifying ultimate (i.e. evolutionary) mechanisms that are likely to have supported the evolution of cooperation in humans.

Moreover, this research has also informed and constrained research into the cognitive and motivational mechanisms that proximally support cooperation. For example, researchers (Rusch & Luetge, 2016; Tomasello et al., 2012) have begun to investigate the psychological implications of a different theory about the evolutionary origins of human cooperation, namely Roberts' (2005) 'interdependence hypothesis.' According to Roberts' theory, humans' tendency to cooperate arose evolutionarily in a period in which our ancestors lived in small groups of individuals whose interests were largely interdependent, and for whom it was therefore not typically beneficial to act selfishly to the detriment of other group members.

We reasoned that if Roberts' (2005) interdependence hypothesis is correct, then repeated coordination with the same partner should boost cooperation independently of any effects of coordination upon trust. This is because repeated coordination may serve as a cue to participants that their partner is interdependent with them, engendering a sense of commitment to that partner. To test this, we designed and implemented a paradigm in which participants performed a repeated joint decision-making task with the same partner in one block (Partner Condition), and with a different partner on each trial in a separate block (Stranger Condition). When both players coordinated on the same option, both were rewarded. On some trials, participants were offered outside options presenting varying degrees of temptation to defect. Participants were informed that they were coordinating with a partner or partners who were in the lab with them. The results of their choices, as well as the trajectories of their mouse movements, indicated that participants in the Partner Condition were more resistant to the temptation to defect. These results support the hypothesis that repeated coordination with a partner can elicit a sense of commitment, leading people to resist tempting alternatives and thereby sustaining cooperation through fluctuations in individuals' interests.

Michel, Matthias (Université Paris-Sorbonne)

The Mismeasure of Consciousness

Abstract: As for most measurement procedures in the course of their development, measures of consciousness face the problem of coordination, i.e., the problem of knowing whether a measurement procedure actually measures what it is intended to measure. I focus on the case of the Perceptual Awareness Scale to illustrate how ignoring this problem leads to ambiguous interpretations of subjective reports. In turn, I show that empirical results based on this measurement procedure might be systematically misinterpreted.

Miletić, Tomislav (University of Rijeka)

IEM and artificial embodiment

Abstract: Certain judgments whose expressions involves the use of the pronoun "I" are taken to be immune to error through misidentification relative to the first pronoun (IEM). For instance, my judgment that I have crossed my legs under the table can be mistaken in several ways, however, it cannot be mistaken that is "I" who have crossed the legs, nor that those legs are mine. As Gareth Evans (1982) explained, the self-specific kind of access I have to my body through the body sense of proprioception, touch and balance give rise to judgments with the IEM property. As such, I may be mistaken about the position of my legs under the table, but I cannot be mistaken that those legs are mine since the information appropriately provided by the body senses' privileged access to my biological body neither requires nor is based on the additional identification. Yet, as Evans famously introduced us to, one can imagine that one's sense of proprioception is linked to another's person body and not one's own and in this case, seemingly, knowing that certain legs are crossed, via proprioception, no longer guarantees that the legs which are crossed are truly mine as the causal link securing proprioception with my biological body is no longer secured. In this paper, I investigate the significance recent embodiment experiments and empirical findings have for this account of IEM. I argue that although the present technological limitations do not permit one to make self-ascriptions of the artificial body's bodily properties in the required appropriated manner to secure their IEM status it does reveal how it is possible, in principle, to do so. As such, bodily self-ascriptions grounded in bodily experiences of artificial bodies can give rise to judgments with the IEM property.

Munneke, Gert-Jan (University of Amsterdam),Fernando Raymundo Velazquez-Quesada (University of Amsterdam), Jakub Szymanik (University of Amsterdam) and Sonja Smets (University of Amsterdam)

Moral Preferences and Dynamic Awareness: A Formal and Empirical Investigation

Abstract: How do people generate and update their moral judgments in an environment where available information changes? To shed light on this question, we first provide a logical formalisation of MFT, proposing a computational mechanism that can account for moral reasoning and consequently empirically test its predictions with a behavioral experiment. The moral foundations' theory (MFT) centers around five moral principles (do not Harm (H), act in a Fair (F) way, be loyal to your Group (G), favour Purity (P), respect legitimate Authority (A) and traditions). Given two of these principles, while some people might find them equally important, others can find one strictly more important than the other. Thus, each person can create an ordering (more precisely, a total preorder). Among these principles, and there are 541 ways to do this (the number of total preorders in a domain with 5 elements; alternatively, the 6th element of the Fubini numbers sequence). However, some of these orderings are 'more common' than others; in fact, Haidt, Graham and Joseph (2009) report that people (living in the United States) can be roughly classified into four different classes (orderings):

secular liberals: {H,F}>A>G>P;
libertarians: F>H>A>G>P;
religious left: H>F>A>G>P;
social conservatives: A>G>H>{F,P}.

If the judgement of a given situation boils down to a conflict between some of these principles, people will tend to follow their orderings. Of course, this will be the case when they are aware that the judgement indeed boils down to a conflict between principles. MFT is a verbal psychological theory. We offer a computational cognitive model based on the dynamic epistemic logic inspired by MFT. The model proposes an explicit computational and logical mechanism for moral reasoning. The model, called awareness

preference model (see Fagin and Halpern 1987, Dégremont and Kurzen 2008), represents not only the ordering an agent might have among situations in which each principle might or might not hold (i.e., orderings over sets of principles) but also the principles the agent is currently entertaining (i.e., the ones she is currently aware of). Thus, the model can predict what would be an agent's response to a moral situation, according to the agent's preferences over moral principles and her current awareness. The model can also predict and explain moral dilemmas as situations in which epistemic updates lead to circular moral preferences. We test the model's predictions in the following way: we have devised two stories whose initial judgement depends on whether a given principle is cherished, and then enrich them (twice in each story) by making the participant aware of another principle that is relevant to the judgement. To simplify the work, only three principles are involved in each story, so participants can be arranged in 13 different groups, according to their self-reported preferences over these principles. The experimental results support the model. It correctly predicts the moral choices of subjects faced with a variety of moral dilemmas. We will report the experimental and modelling details in the talk.

Newen, Albert (Ruhr University Bochum)

The embodied self in the predictive mind

Abstract: Do we have to presuppose a self to account for human self-consciousness or not? If so, how should we characterize the self? These questions are discussed in the context of the no-self position hold by Metzinger (2003, 2009) and the claim that the only self we have to presuppose is a narrative self (Dennett 1992; Hardcastle 2008; Schechtmann 2007) which is mainly an abstract entity. In contrast to these theories I argue that we have to presuppose an embodied self while this embodied self is not a metaphysical substance or any entity with stable necessary and together sufficient conditions. Self-consciousness result from an integration of an embodied, basic affective flow with an intentional object (the self as agent or as center of imagination or thought) where this integration remains anchored in an embodied self. This embodied self is a flexible and varying entity which we can only adequately account for with a pattern theory of the self (in line with Gallagher 2013). Furthermore, it is outlined how this pattern theory of the self fits into the predictive coding framework. This also answers the open question whether self-representation is prior to world-representation or the other way around. The principle organization of a mechanism of building up a self-model is such that both types of representations are activated and developed always in parallel. Modelling oneself is an always activated process when interacting with the world like a shadow is present when walking in the sun.

Niemeck, Maik (University of Freiburg)

Personal or Sub-Personal Level Explanation of First-Person Awareness? – The Argument from Language Use Reconsidered

Abstract: Many philosophers – such as Bermúdez (1998), Guillot (2016), Lang (2010), Musholt (2015), Nozick (1981), O'Brien (2007), Recanati (2007, 2013) or Zahavi (1999) – believe that the mastery of the first-person pronoun already presupposes a form of non-linguistic experiential self-awareness. A central argument for this claim usually proceeds as follows.
The linguistic meaning of first-person pronouns is most commonly understood as something along the lines of Kaplan's concept of character, which is for "I": "I" rigidly refers to the user of a specific token of it. (cf. Kaplan 1989: 505) When we think about us using a first-person pronoun (or its cognates), we usually know that we do so and thus that we are the referents of these tokens. This form of self-thinking can be in harsh

contrast to cases in which we use proper names or definite descriptions that happen to refer to ourselves without our knowledge, as Castañeda (1966), Lewis (1979) and Perry (1968) have influentially established. The question that arises from these observations is how we are able to perform reflexive first-personal self-reference. This ability appears to demand more than a mere understanding of the linguistic meaning of the first-person pronoun, since it only specifies that the pronoun refers to the user of the pronoun, whoever that might be in a specific situation. Adherents of the claim that there is phenomenal self-awareness declare that only non-linguistic information given in experience can fill the gap and deliver us with the information that we are currently the thinker of specific I-thoughts. Thus, phenomenal self-awareness is considered to be an experiential condition for reflexive linguistic self-reference.

In my talk, I will firstly demonstrate that most proponents of the presented argument fail to establish their conclusion, even though I believe that the conclusion is true. The argument sketched above only justifies the claim that contextual information is needed in order to identify oneself as the user of an occurrence of the first-person pronoun. However, it does not justify the thesis that this information needs to be given consciously and thus the existence of phenomenal self-consciousness.

In the second part, I shall present an argument for the assumption that some form of phenomenal self-consciousness is needed to reflexively think about oneself via first-person pronouns and that sub-personal processes might not be able to fulfil this function. In order to do so, I will draw on the general discussion about the cognitive role of conscious contents. In particular, I will rely on results recently proposed by Frith/Metzinger (2016) and Frith (2010) which suggest that the specific function of consciousness is to provide information for communication rather than for complex problem solving.

Niemeck, Maik (University of Freiburg)

Demystifying IEM – About the Irrelevance of Immunity to Error through Misidentification for Theories of Self-Awareness

Abstract: Many philosophers – such as Bermúdez (2016, 1998); Musholt (2015, 2011), Recanati (20013) or Zahavi (2004, 1999) – believe that an understanding of the so-called Immunity to Error through Misidentification ("IEM" henceforth) (cf. Shoemaker 1964) of self-ascriptions of currently conscious states is essential to develop adequate accounts of various forms of self-awareness. Hence, a proper understanding of self-awareness is apparently entangled with an understanding of IEM. To give just one recent example for such a view: "[IEM] is an epistemological phenomenon that is prima facie related to the special character of de se thoughts expressed with uses of the first-person. One should hence expect that accounts of de se thoughts explain or otherwise illuminate IEM" (Garcia-Carpintero 2015: 2).

Contrary to these positions, I will attempt to show that IEM is in no regards relevant for theories of self-awareness. My argument for this claim proceeds as follows. At first, I will lay out why IEM is not to be understood as some form of reference failure and why it is most commonly regarded as a relational property of a thought and its epistemic base. Thereafter, I will maintain that if this is true then IEM is not to be understood as a weakened form of infallibility, as many philosophers have explicitly claimed or at least ambiguously hinted at. Moreover, I will demonstrate that a definition of IEM based on this assumption either is not able to restrict the range of cases to which the phenomenon applies – which would be needed to point out any relevance for De Se awareness – or that it actually is a quite trivial observation not worth mentioning. In conclusion, even when we assume that there are thoughts which are IEM – i.e. thoughts with an identification free epistemic base – it is a too widespread phenomenon to be relevant for theories of self-awareness.

In the second part of my talk, I will proceed by pointing out another way to make sense of one central intuition behind what philosophers call IEM. Many philosophers think that IEM amounts to a weakened form of infallibility. They maintain that specific thoughts cannot be wrong in a certain respect although they might be wrong tout court. I will show that one central idea behind this is the assumption of a necessary link between introspective phenomenal consciousness and property-possession. I will propose that there might be subject-centered sources of evidence that are necessarily about the individual that accesses them. I will call a definition of IEM based on this assumption "metaphysical IEM". Something very close has been suggested by Seeger (2015). I will argue that the definition of metaphysical IEM is able to restrict the range of cases to which IEM applies and moreover that it is not trivial to find instances of that definition. However, I will proclaim that if this is what IEM is really about then it actually has nothing to do with an immunity to misidentification and because of that is irrelevant for theories of self-awareness.

O'Conaill, Donnchadh (University of Fribourg)

Subjectivity and Non-Objectifying Awareness

Abstract: A long tradition holds that each subject is necessarily aware of her experiences as she has them (I term this the subject's inner awareness). Inner awareness is often said to constitute the subjectivity of our experiences, their distinctively first-personal character (Levine 2001; Kriegel 2009).
There are different conceptions of inner awareness, e.g., as perceptual in nature or as a kind of thinking. An important line of thought found in phenomenological thinkers such as Husserl and Sartre holds that inner awareness is non-objectifying: "when we are absorbed or immersed in our daily concerns and simply live through our experiences, they are not given as objects; they are not something we observe from a distance and they do not stand opposite us" (Zahavi 2005, 64).
The idea of non-objectifying awareness is often criticised as obscure (Kriegel 2009, 101-106; Schear 2009, 103-104; Hoerl 2013, 388). What is needed is a positive account of what is it for awareness to be non-objectifying. To provide such an account, I first address what it is to be aware of an entity, x, as an object. It is difficult to provide necessary conditions, but the following seem sufficient: being aware of x as bounded and demarcated from other entities, and as a particular instance of a general kind. To be aware of x but not in this way would be, in an important sense, to have non-objectifying awareness.
Next, I appeal to P.F. Strawson's work on feature-placing statements (1959). These are statements where the distinction between universals and particulars does not apply. In these statements a difference cannot be drawn between different particulars instantiating the same universal at different times, and the same particular instantiating this universal over time. Examples of feature-placing statements are 'Now it is snowing' or 'There is water here'. In contrast, a non-feature-placing statement such as 'This is a cat' involves a distinction between a universal term ('cat') and the particular instance of it picked out in the statement.
I suggest that inner awareness is a feature-placing form of awareness: it does not involve a distinction between universals and particulars. In inner awareness one is aware of one's experiences, which are particular events or episodes, but one is not aware of them as particulars. For example, the content of one's inner awareness might be 'Now it is painful', as opposed to 'This is an experience of pain'.
This allows us to understanding how inner awareness can be non-objectifying. In inner awareness experiences are not presented as particular instances of a general kind, nor as demarcated from other experiences of the same kind. One is not aware of discrete experiences, but of a varying stream of experiencing. In reflection, one can pick out particular experiences from the stream, considering them as bounded individual events with a specific character and temporal extension. This would be to consider

them as objects, albeit objects to which one stands in a specific relation. But in inner awareness they are not given in this way.

Ochs, Jordan (University of Connecticut) and Dorit Bar-On (University of Connecticut)

Speaking Your Mind in Your Mind: Inner Speech and Self-knowledge

Abstract: Our knowledge of our own present states of mind is thought to be both distinctive and privileged. Gilbert Ryle has famously offered a deflationary account of self-knowledge, arguing that our inner speech episodes (ISEs) could serve as a privileged (and possibly also distinctive) evidential basis for self-knowledge of mental states. Ryle's account has been, for the most part, rejected. However, several authors have recently attempted to revive the Rylean account by way of explaining the role of inner speech in self-knowledge.In a paper that has just appeared, we argue against two of these Neo-Rylean accounts and identify desiderata for a satisfactory theory of inner speech as it features in self-knowledge of occurrent mental states. Here, we develop more fully a positive account that can meet these desiderata. We start by recapping the two neo-Rylean accounts we reject and rehearse the extracted desiderata. We then turn to our main proposal, which is that the neo-expressivist view of avowals has the right overall shape to satisfy the desiderata we have laid down. We then explain how the neo-expressivist account as applied to inner speech can accommodate the idea that we have distinctive and privileged knowledge of our own thoughts.

Orban, Krisztina (University of Tuebingen) and Hong Yu Wong (University of Tuebingen)

Multimodal Bodily Immunity to Error through Misidentification

Abstract: On a classical account of bodily immunity to error through misidentification (IEM), bodily self-ascriptions are immune when they are based solely on perception of one's own bodily properties 'from the inside' (Evans 1982). De Vignemont (2012) has criticised this account on the basis of the multimodal character of internal perception. She proposes a multimodal account of bodily IEM based on an analysis of cases of visual and multimodal experiences that are supposed to ground immune bodily self-ascriptions. We argue that de Vignemont's account of multimodal bodily IEM is either too narrow or too board and that there is a revised version of the internal view which allows for multimodal bodily IEM.

Ozols, Davis (University of Fribourg), Pascal Gygax (University of Fribourg), Richard Breheny (University College London) and Didier Maillat (University of Fribourg)

The effect of previous beliefs and argumentative engagement on majority agreement

Abstract: Argumentation is inherently linked with persuasion and in most everyday cases it is used as a tool to justify ones beliefs or actions and to convince others of their correctness or necessity. Due to the social nature of our everyday lives we are exposed to argumentation of various contents and from various sources on a daily basis. This means that people must possess fast and frugal ways of analysing argumentation to make this form of communication effective and to potentially avoid misinformation and manipulation. A mechanism that could serve this role - Epistemic Vigilance - has been proposed by Sperber et al. (2010). Epistemic Vigilance provides two main routes for vigilance towards communication - content and source. This prompts one to investigate a question: in what way can such a mechanism be biased and in which cases it might produce suboptimal results?

In this paper, we investigate biases that can affect peoples' vigilance towards communicated information and review cases where epistemic vigilance might fail to safeguard against subpar arguments because of said biases. For this, we look at the persuasiveness of argumentative and linguistic patterns referring to majority agreement, known in argumentation literature as arguments from popular opinion – or ad populum arguments. In our experiments, we specifically look at the influence of (i) preexisting beliefs and (ii) willingness to defend them on the way people evaluate arguments with reference to group agreement, with the idea of assessing whether the myside bias (i.e. the tendency to select and evaluate information in favour of preexisting beliefs (Stanovich et al., 2013)) affects their evaluation.

Previous empirical investigations into the ad populum phenomenon are rather scarce, with few notable exceptions in the Bayesian paradigm of argument evaluation and reasoning (Hornikx, 2005; Hahn and Oaksford, 2007; Jönsson et al., 2015; Hahn and Hornikx, 2016). Furthermore, experimental studies, which have investigated the role of majority opinions in argumentation usually, have relied on items presenting opinions with a clear correct answer (Moore, 1921; Harris and Hahn, 2009; Jönsson et al., 2015) or in combination with expert opinions (Hornikx et al., 2018). Our approach differs in that (i) we investigate non-expert sources and (ii) we use non- normative items (i.e., claims relating to deontic, ought-to statements), which are thus devoid of clear-cut correct answers. In this study, we hypothesise that biases play a role in the effectiveness of such arguments, and specifically the myside bias, as its influence in reasoning tasks has already been established (Wolfe, 2012). We connect our results to discussions on argumentative fallacies, myside bias (Stanovich et al, 2015) and the argumentative theory of reasoning (Mercier and Sperber 2011, 2017).

Perak, Benedikt (University of Rijeka)

Using the corpus analysis in analyzing Cultural Models and linguistic activation of Affective Experience

Abstract: The paper deals with the identification, extraction and cross-linguistic comparison of the psychological concepts subcategorized as perceptual, affective and cognitive concepts, and their relationship with material and socio-cultural concepts. What are the salient conceptual structures of the linguistic construals that are used to express psychological perceptions, affects and cognitions as entities (nouns), processes (verbs) and properties (adjectives, adverbs) in the communication? The methodology of this ontological corpus-based study includes three phases. The first phase deals with the construction of the Ontological Model of Concepts and Linguistic Constructions database that aims to formalize the meta-data about the ontological features of the psychological concepts and their relation with material and socio-cultural concepts. The ontological model is theoretically grounded in the system theory (Emmeche et al. 1997, Baas & Emmeche 1997, El-Hani & Emmeche 2000, Searle 2006, Capra & Luisi 2014), and cognitive approaches to the categorization (Rosch 2005). The ontological model is stored in a graph property database Neo4j (https://neo4j.com/). The second phase includes the extraction of the nominal, adjectival and processual lexical concepts related to the psychological phenomena from the large corpora of Croatian (hrWaC 2.2) and English (enTenTen13) using the SketchEngine API and Reldi tokenizer and parser. The psychological domains are extracted using the syntactic methods of paradigmatic similarity score for the co-occurrences in the coordinated construction [x and y] for nominal lexemes Sketchengine platform (https://the.sketchengine.co.uk). Using graph algorithms for community detection the lexemes in the coordinated linguistic constructions are classified for their syntactic-semantic domains.
The third phase examines the ontological status of the lexemes and superimposes logical inferences on the semantic-syntactic constructional relations of the language-specific knowledge.

This empirical approach sets the dynamic systems theory as the epistemological basis for studying ontological questions of the syntactic-semantic relations expressed in language, its metaphoricity, dynamic network relationships, nonlinearity, emergence, complexity, hierarchy, ontological contingency and congruence of the conceptual organization of psychological concepts (Larsen-Freeman 2015).

References:
Emmeche, C., Køppe, S., & Stjernfelt, F. (1997). Explaining emergence: towards an ontology of levels. Journal for general philosophy of science, 28(1), 83-117.
Baas, N. A., & Emmeche, C. (1997). On emergence and explanation. Intellectica, 25(2), 67-83.
Capra, F., & Luisi, P. L. (2014). The systems view of life: A unifying vision. Cambridge University Press.
Larsen-Freeman, D. (2015). Ten 'lessons' from complex dynamic systems theory: What is on offer. Motivational dynamics in language learning, 11-19.
Rosch, E. (2005). Principles of categorization. Etnolingwistyka. Problemy języka i kultury, (17), 11-35.

Petrolini, Valentina (University of Cincinnati) and Marco Viola (Moscow State Pedagogical University)

From Affect Circumplex to Affect Pendulum? Insights from psychopathology on the relation between arousal and valence

Abstract: In the past thirty years, the debate on emotions and other affective states has flourished in a number of sub-fields within psychology and philosophy. Notably, there have been several attempts to construct models that would successfully describe affective experience by individuating its basic components. Pioneering work in this sense has been conducted by James Russell (2003). On his account, our affective lives are characterized by a baseline neurophysiological state called Core Affect, which consists in a certain degree of arousal (activation) combined with a certain degree of valence (pleasure-displeasure). These two allegedly independent dimensions are usually represented in what Russell dubbed the Affect Circumplex (1980). For example, an emotion such as anger would be characterized by high arousal and negative valence, whereas a calm mood would be characterized by low arousal and positive valence: these two states would thus occupy opposite corners in the Affect Circumplex.

In this talk we draw on Russell's notion of Core Affect to explain some salient features of psychopathology, with a focus on depressive disorders. The argument is divided into three parts. In §1 we argue that, despite the current lack of knowledge about underlying mechanisms, Russell's Affect Circumplex is a legitimate and well-suited tool for describing the affective states underlying a number of mental disorders. In §2 we focus on a recent study conducted by Russell and his collaborators (Kuppens et al. 2013) in order to raise a challenge to their account. In brief, we argue that some important features of psychopathology concerning the relation between arousal and valence can be captured only if affective states are measured diachronically (i.e. over time) as opposed to synchronically (i.e. at a given time t). Indeed, we believe that the almost exclusive focus on synchronicity has obscured some important interactions between the two dimensions of Core Affect. In §3 we sketch a working hypothesis about a currently unexplored interaction between arousal and valence, that we dub the "Core Pendulum Hypothesis". Specifically, we suggest that – in ordinary as well as pathological cases – arousal could have a modulating effect on valence, acting like the rod in a pendulum: the longer the rod, the greater the oscillation. When arousal increases, valence tends to oscillate more: this may happen both in non-pathological situations (e.g. when one is over-caffeinated) and in pathological ones – e.g. bipolar depression. On the contrary, when arousal decreases it becomes easier to get "stuck" in a particular valence: this happens whenever we have trouble snapping out of a mood, and

more severely in cases such as unipolar depression. Such hypothesis may elegantly account for the weak but consistent V-shaped relation between arousal and valence found by Kuppens et al. (2013). We conclude the talk by devising some possible ways in which the Core Pendulum Hypothesis could be tested.

Prochownik, Karolina (Ruhr University Bochum) and Matthias Unterhuber (Ruhr University Bochum)

Blame blocking in legal experts? Experimental evaluation of the (legal) expertise defense

Abstract: Imagine two agents trying to commit a crime. One agent intentionally tries to kill a person, and he undertakes all the necessary steps to do so. Unfortunately, the harm does not occur due to unforeseen circumstances, and the victim is unharmed. Another agent wants and does the same as the first one; his actions are equally unsuccessful in bringing about the harm. However, the victim dies completely independently of his actions and plans.

Previous psychological findings have shown that in this kind of cases lay people assign significantly less blame and punishment to the second agent (who tries to commit a crime and fails, but the intended harmful outcome happens independently of his actions) than to the first one (who attempts to commit a crime and fails). This puzzling effect was named "blame blocking" (Cushman, 2008).

However, contrary to the lay people`s judgments, the law treats these two cases of failed attempts likewise. Since the perpetrators had the same intentions and performed the same set of actions, they should receive the same or at least similar amount of punishment. Circumstantial features of the situation such as the occurrence of an independent causal chain leading to harm should be irrelevant for assessing criminal liability for the perpetrator´s attempt. Consequently, we would expect lawyers or people with legal expertise to evaluate such cases according to the letter of the law, and thus to be resistant to the "blame blocking effect". Nonetheless, many of the recent findings in experimental philosophy undermine the picture that experts are in any way less susceptible to biases that characterize non-experts´ evaluations of hypothetical scenarios.

In the paper, we will present a set of experiments on the blame blocking effect conducted on people with and without legal expertise in Poland. We successfully replicated the blame blocking effect with lay people. Regarding the legal experts, we found a significant difference in the susceptibility to the blame blocking effect, depending on the number of years of legal studies completed. "Beginning legal experts" (people who studied law for less than one year) tended to still exhibit the blame blocking. By contrast, "advanced legal experts" (people who studied law for more than one year) were not susceptible to the blame blocking effect. These results suggest that legal expertise may have a positive impact on attenuating (at least some) psychological effects and cognitive biases in reasoning and decision making.

We discuss the implications of these findings for the legal system and the previous research on expertise defense in general, and expertise defense for legal professionals in particular (Kneer & Bourgeois-Gironde, 2017).

References
Cushman, F. (2008). Crime and punishment: distinguishing the roles of causal and intentional analyses in moral judgment. Cognition, 108(2), 353-380.
Kneer, M. & Bourgeois-Gironde, S. (2017). Mens rea ascription, expertise and outcome effects. Professional judges surveyed. Cognition, 169, 139–146.

Quinon, Paula (Lund University) and Peter Gärdenfors (Lund University)

One-to-one correspondence between temporal and spatial representations of numbers in the development of counting

Abstract: We present a model of the number concept acquisition. When children learn to count, they need to grasp, or understand, several partial concepts. In consequence, as we argue, several cognitive capacities are involved. We look into the conceptual structure of the number concept and of the concept of counting, and on this basis we present a model how counting is learned. We distinguish three components of counting: (1) understanding that collections can have properties; (2) learning the sequence of numerals; and (3) learning how to establish a one-to-one mapping from numerals to collections. The third component has not been analysed in the context of number acquisition.

(1) Collections can have different properties. Some of these properties are closely tight with the physicality of objects forming collections, for example weight and location: "These beans weigh 500 grams". "The radishes are in the plastic bowl in the fridge". Many properties are, however, more abstract collections can be ordered or unordered, uniform (consisting of the same type of objects) or mixed, dense or spread out. In particular, collections have numerosities..

(2) Numerals are a crucial invention in structuring the way we experience the world. We have a rough sense of quantity without them, but numerals are necessary to impose exact, discrete quantities in our experiences. At first, numerals are learnt as a list of arbitrary words – up to ten or twelve depending on the child's language. During the child's development, the list becomes more systematic and most languages have a method for recursively generating new number words.

(3) The key skill in learning to count is the ability to establish a one-to-one mapping between the list of numerals and the elements of a collection. First, we observe that a child needs to grasp several other mappings: a one-to-one mapping between individual elements from different collections, or a mapping between different cardinalities and numerals. This leads to grasping of different types of successor (e.g., adding one item to a collection corresponds to moving of one element on numeral line) and to understanding that the last numeral uttered when the mapping between elements and numerals finishes is the numerosity of the collection.

Second, we observe an important aspect that has been neglected in earlier literature on counting. The collection of elements that are counted have a distribution in space (e.g., the cookies to be counted are located on a plate), while the recital of numerals and the pointing to the objects is a process that is extended in time. The counting involves a mapping from the time domain to the space domain. Thus counting involves a form of cross-modality.

In our presentation we propose a model that accounts for the mapping between the time and the space domains. We also propose an in-depth analysis of the temporal and the spatial dimension of counting. For instance, we observe that counting is typically done in a rhythmical manner. The equidistance in time of rhythmical counting mimics the equidistance in numbers.

Rajnović, Martina (University of Rijeka)

Do psychopaths lack empathy?

Abstract: Psychopathy is a personality disorder characterized by grandiose sense of self-worth, deception, manipulation, dominance, superficiality, shallow affect, lack of empathy and irresponsible and impulsive socially deviant (not necessarily criminal) behavior (Hare, Neumann 2009). Recently it has attracted the interest of philosophers as a paradigm case for testing philosophical accounts of moral capacities.

One area of the debate is the significance of psychopathy in the investigation of the psychological capacities that underlie moral understanding. Some philosophers have argued, for instance, that some empathic impairments in psychopaths explain the impairments in their moral understanding (Prinz 2006, Aaltola 2014, Nichols 2004).

Evaluating these arguments should involve investigating at least three issues: is there good evidence for concluding that psychopaths lack moral understanding, what is the role of empathy in moral understanding, and do psychopaths lack empathy?

I will focus on the last issue. I argue that these arguments appear to rely on an uncritical extrapolation of philosophical insights from psychological and neuropsychological research. Once we devote close attention to issues concerning the definition of empathy and its measurement in experimental settings, the empirical evidence for the claim that psychopaths lack empathy is not so compelling.

Despite a long tradition in philosophy and behavioral psychology research, empathy has no universally accepted definition. Furthermore, it is often not distinguished enough from emotions like sympathy, emotional contagion, personal distress and attitudes like perspective taking. In general, empathy is taken to be an emotion a person feels when she is happy, sad or embarrassed for someone else because of that other person's (presumable) feelings in a given situation. (Maibom 2009, Davis 1994, Hoffman 2000, Sober and Willson 1998). However, we must be aware that there are two kinds of empathy: affective and cognitive. Affective empathy is usually taken to be the drive to respond appropriately to another's emotions. This kind of empathy happens automatically, and often unconsciously. Cognitive empathy is the largely conscious drive to understand the state or feelings of someone in a given situation by imagining how it would feel if we were in her shoes. It is often seen as precursor to affective empathy.

There are at least two types of measures used to establish empathy levels: situational and dispositional measuring. Situational measuring is useful for discovering what kind of situations usually give rise to empathy and what kind of reactions empathy motivates (Maibom, 2014). Dispositional measuring, which seems more interesting for present purpose, measures the tendency of an individual to respond empathically to others.

Empirical studies concerning psychopathy rely on several methods for measuring empathy, I will focus on the following methods often mentioned by philosophers: startle reflex (Patrick et al. 1994), electrodermal response (Lykken 1957, Hare 1978, Ogloff & Wong 1990, Fowles 2000, Flor et al. 2002), heart rate variability (Eisenberg et al., 1991), self-reports and questionnaires (Hogan 1969, Mehrabian & Epstein 1972, Davis 1980, 1983, and 1994). I will argue that the empirical studies based on these methods cannot unequivocally support the conclusion that psychopaths lack empathy.

Roessler, Johannes (University of Warwick)

Reflective perception and the subtraction argument

Abstract: Perceiving objects around us not only enables us to know what objects are like but also to understand how we know what we know. For example, if you can visually tell that the thing before you is a lemon, you will also be able to reflect that you can see it's a lemon (and perhaps offer at least a rough indication of how you can tell). According to some philosophers, the capacity for this sort of reflective understanding of one's perceptual knowledge is an inseparable aspect of the capacity for perceptual knowledge itself. (McDowell 2011 is one example.) In this talk I consider, and try to disarm, an influential

developmental objection to this view, which I will call the subtraction argument. The argument goes like this:

(1) Young children have the capacity to acquire perceptual knowledge of objects around them.
(2) Young children lack the capacity to reflect on the source of their perceptual knowledge.
(3) So the human capacity for perceptual knowledge is not inherently reflective.

The diagnosis I want to explore says that the subtraction argument is guilty of equivocation. What is true is that children have a *rudimentary* capacity to acquire perceptual knowledge, and that they lack a *fully developed* understanding of the source of their knowledge. Read in this way, though, (1) and (2) lend no support to (3). The argument only goes through if the premises are read as saying that children have a fully developed capacity for perceptual knowledge, and lack even a rudimentary grasp of the source of such knowledge. On that reading, both premises are wrong.

I think the two main questions that defenders of this line of response need to address are these: (a) in what sense is young children's capacity for perceptual knowledge limited or rudimentary? (b) How is that sense related to well-documented limitations in their reflective understanding of the source of perceptual knowledge? The suggestion I will pursue is that the answer to both issues turns on the nature of perceptual judgement. Using perception to gain propositional knowledge is not the same as perception causing certain kinds of subpersonal representations. It involves being able to answer questions, i.e. to make perceptually grounded judgements or assertions. Arguably, *full* mastery of that ability requires appreciating that judgements are open to certain sorts of challenges, notably the question 'How do you know?', which can demand a reason for accepting your claim to knowledge. Children's difficulties with the 'perspective problem' (in the sense of Perner et al 2003) created by that question may help to account for their poor performance on certain tasks probing their understanding of the source of knowledge. Yet, even early in development children seem to have both some understanding of the enabling condition of perceptual knowledge and some facility for expressing and sharing such knowledge, as shown e.g. by the use of 'holophrases' and 'proto-declarative' pointing (Moll 2012). Consideration of the developmental evidence, I argue, supports a nuanced reading of (1) and (2).

Sanhueza, Sebastian (Catholic University of Maule)

Heracliteanism about Experience

Abstract: Aiming to show how questions of perceptual ontology bear on questions of perceptual phenomenology, Matt Soteriou conceives perceptual experiences as phenomenally conscious states—i.e. temporal items which obtain at a time or continue obtaining over time—that constitutively depend on processes—i.e. temporal items that unfold over time—of a special phenomenally conscious kind. Heavily drawing on Brian O'Shaughnessy's work, Soteriou's stance is to an important extent Heraclitean, for the relevant experiential processes are supposed to be in permanent or necessary flux. That is, while ordinary physical or psychological but non-experiential processes typically involve continuous changes from state to state, experiential processes would allow for no stative analysis: according to Soteriou, phenomenally conscious processes are processive, occurrent, or dynamic down to their quasi-infinitesimal parts. In the present talk, I aim to put pressure on this Heraclitean ontology of experience. In a nutshell, the driving motivation of my negative assessment concerns the fact that, by drawing on necessarily dynamic processes, Heracliteanism about experience obscures how perceptual experiences relate to the natural world.

The present task is divided into three parts. First, I outline Soteriou's view of perceptual experiences as mental states constitutively dependent on necessarily dynamic processes. Secondly, I flag two difficulties posed by the processive items thus invoked: on the one hand, they lack clear conditions of individuation; and, on the other, they actually obscure the sense in which perceptual experiences are dynamic. Finally, I go on to argue against Soteriou's main motivation for introducing such entities into our ontological framework, namely, the thought that an ontology of necessarily dynamic processes elegantly accommodates the phenomenology of temporal passage. In a nutshell, I argue that, even if we concede the controversial phenomenological datum from which Soteriou's defense takes off, it is by no means clear that experiential Heracliteanism has the aforementioned explanatory power.

Sarıhan, Işık (Central European University)

Experiential Acquaintance with Abstract Entities

Abstract: This paper postulates and explores an epistemic state which is inherent to all phenomenal experiences: Direct acquaintance with possibilities. Relying on a representationalist framework and a common-factor theory of perception and hallucination, I argue that having a perceptual experience involves an acquaintance with abstract entities. These abstract entities are possibilities, for instance the possibility of a yellow sphere existing in space and time, and this epistemic state is made possible by, indeed necessarily follows from, the experiential non-conceptually representation of a yellow sphere. Experiences, therefore, bestow knowledge upon subjects about sensible qualities, such as the colour yellow, even in hallucinatory or illusory cases, and this epistemic state is separate from, and not in conflict with, the direct awareness of objects and their properties in veridical experience. With the help of several thought experiments, I demonstrate that pure cognition lacks this element. I argue that puzzles of perceptual experience are better solved by postulating this epistemic relation, and erroneous arguments like the Argument from Hallucination result from a misunderstanding of this epistemic feature of sensory experience, leading to the postulation of sense-data and qualia. I further apply this theory of acquaintance to the philosophical puzzle regarding Mary the Colour Scientist, providing an explanation, without postulating qualia, of why she lacks some knowledge in her pre-experiential, purely cognitive state.

Sarkia, Matti (University of Helsinki)

Mechanistic Explanation as an Interdisciplinary Endeavor: the Case of Shared Task Representations

Abstract: The phenomenon of joint action has become a topic of convergent interest in cognitive science (Knoblich et al. 2010), developmental psychology (Rakoczy 2017) and evolutionary anthropology (Tomasello et al. 2005; Tomasello 2014) in recent years. Roughly speaking, joint action occurs when two or more individuals coordinate their behaviors over space and over time so as to realize a common goal (Knoblich et al. 2010). A rough-and-ready division of labor between these disciplines can be delineated as follows: Evolutionary anthropologists seek to understand what forms of joint action and social coordination set apart the human species from our nearest primate relatives. Developmental psychologists seek to understand how the capacity for joint action develops during human ontogeny. And cognitive scientists seek to find out what psychological mechanisms bring about joint action. While satisfactory as a first sketch, this characterization leaves open a host of important details about how these disciplinary perspectives on joint action latch together, where are the joints at which they interlock, and what are the nuts and bolts that tie them together into an overall coherent story. To answer such questions, a more holistic methodological perspective is called for.

This paper draws on the literature on mechanistic explanation in the philosophy of science (Bechtel&Richardson 2010; Craver&Darden 2013) so as to provide a comprehensive account of how approaches to joint action in cognitive science, developmental psychology and evolutionary anthropology may complement one another by filling in different details to a mechanism scheme describing the stable dispositional features and typical causes of joint action. To illustrate my view, I will discuss recent research on shared task representations as a possible mechanism for joint action. Shared task representations were originally postulated to explain observed behavior in the experimental paradigm described as the Joint Simon Task, which involves two individuals performing a similar task alongside one another (Sebanz et al. 2003). Dolk&Prinz (2017) have recently argued that current evidence for shared task representations is inadequate, and that the experimental data that has been brought up to support their existence can be accounted for by their more parsimonious proposal in terms of referential coding. After assessing the severity of the challenge posed by Dolk&Prinz (2017), I will argue that the existence of shared task representations would support Michael Tomasello's (2005; 2014) research hypothesis according to which a phylogenetically unique capacity for joint action and shared intentionality played an important role in setting the human species apart on its unique evolutionary pathway. By contrast, the rejection of shared task representations would undermine the tenability of Tomasello's hypothesis, while also leaving a range of experimental data in developmental psychology and evolutionary anthropology partially unaccounted for.

Schulte, Peter (Bielefeld University)

Smelling How the World Is: Do Olfactory Experiences Represent?

Abstract: Our sense of smell has long been neglected by philosophers, despite the fact that olfactory experiences give rise to many interesting philosophical questions. Arguably the most fundamental question is this: are olfactory experiences representational states? In this paper, I will argue for an affirmative answer, i.e. for representationalism about olfaction (or representationalism, for short). First, I aim to show that an argument for representationalism that figures prominently in the literature on olfaction (Batty 2010a, 2010b) lacks dialectical force. I then go on to develop two novel arguments for the representationalist position: (i) the argument from illusion and hallucination, which is based on how experts characterize dysfunctional olfactory experiences, and (ii) the contrast argument, which draws on information-processing considerations. From these arguments, I conclude that representationalism about olfaction is correct: olfactory experiences are (exteroceptive) representations. In other words, when have an olfactory experience, we smell how the world is.

Schuster, Annika (Heinrich Heine University Dusseldorf) and Leda Berio (Heinrich Heine University Dusseldorf)

Mental representations in dialogue: explaining (mis)communication with mentalized frames

Abstract: Research on the structure of mental representations rarely focuses on how concepts are used and modified during communication. However, informational asymmetries and idiolects characterise everyday (mis)communication, revealing differences between the mental representations of participants in a conversation; taking this into account, we argue that the success of meaning coordination in communicative exchanges is frequently determined by making adequate, often implicit, assumptions about each other's conceptual structures and that this is reflected in an appropriate tuning of our own conceptual

representations for communication purposes. We thus focus on those cases of communication whose success is at least partially determined by speaker's and hearer's adjustments to their own conceptual repertoire.

Our aim, then, is to model how individual mental representations are tuned during (un)successful conversational exchanges, focusing on the incorporation of information regarding speaker's and hearer's respective knowledge.
Two fundamental tenets of our model are:

1. Weighted recursive attribute-value structures (frames), which are a cognitively plausible way to explicate the structure of conceptual representations (Barsalou 1999, 2005). This representation format is especially powerful as it is compatible with the main cognitive theories of conceptual structure and it allows the incorporation of probabilistic information.
2. Mentalizing (Theory of Mind) abilities operating on frames to identify communicable content, incorporating contextual information and background knowledge and considering both speaker's and hearer's knowledge. Formally, we treat mentalizing as an operation of adjusting conceptual structures based on second order probability considerations.

Furthermore, to model background shared information, we rely on a distinction between General Common Ground and Immediate Common Ground that incorporates contextual information and background knowledge (Berio et al. 2017) and propose a formalisation. In this way, the model accounts for the role of contextual information in communicative exchanges and for how context interacts with speaker's and hearer's estimation of each other's knowledge.

We will argue that our a model can produce empirically testable predictions regarding the extent to which speakers are able to estimate hearer's state of knowledge and how this impacts communication. One of the predictions of the presented model is that communication is especially facilitated with prototype concepts (Schurz, 2012), where the overlap between speaker's and hearer's knowledge is high by default and less mentalizing is needed; on the other hand, cases where idiolects and technical language are used are more prone to error, given the bigger amount of mentalizing required. We will describe the empirical paradigm that we intend to use to test these predictions; furthermore, we will present some examples of systematic communication failures in doctor-patient communication, which is a field especially prone to communication errors due to informational asymmetries (Musso, 2012), and show how these cases can be fruitfully modelled in our framework.

The model combines a view of mental representations as complex attribute-value structures and an account of the role their online tuning plays in communication.

Serpico, Davide (University of Leeds)

Psychological Traits and Genetic Causation: Guidelines for Defining Phenotypic Traits

Abstract: A major issue in understanding the genetic bases of human psychological traits lies in how hard it is to provide reliable definitions of psychological traits themselves. Characteristics such as intelligence and mental disorders are often poorly defined and, thus, display much complexity as well as fuzzy boundaries. Nevertheless, behavior geneticists understand these characteristics as phenotypic traits and assume they can be analyzed by genetic methodologies generally adopted for the study of simpler biological

characteristics (e.g., stature and pigmentation). For instance, heritability and genome-wide association studies have been widely adopted in the analysis of human psychology. However, the relationship between genes and psychological traits is still unclear and important theoretical controversies afflict behavior genetics. In this paper, I address these controversies by asking how and to what extent genetic causation can spread across higher levels of the biological organization, from molecular to psychological characteristics. This analysis can provide guidelines for identifying phenotypic traits at different levels of organization and will represent an important advancement for philosophy of life sciences, genetics, and psychiatry.

Shardlow, Jack (University of Warwick)

To Experience the Passage of Time

Abstract: In the past two decades there have been a plethora of accounts seeking to explain why time seems to pass. Explanatory accounts have been proposed that appeal to (episodic) recollection (Mellor 1998); to perceiving motion/change (Paul 2010); or to perceiving objects as enduring (Prosser 2012 and 2016). In addition to these accounts, there have been arguments against the claim that we perceive time passing (Hoerl 2014).

In this paper I develop an account of how time seems to pass which has so far been neglected in these discussions. On the proposal I develop, time seems to pass – for experiencing subjects – because of a proprietary phenomenology of agentive wakeful consciousness, rather than supposing that the passage of time is presented in perceptual experience. I suggest that we are independently motivated to endorse this account when we reflect on how subjects experience their own intentional actions.

Smokrović, Nenad (University of Rijeka)

Formal and informal logic: their relationship in reasoning

Abstract: The paper focuses on is the normativity of everyday reasoning (arguably performed in a dialogical argumentative setting) and, more precisely, the relationship between formal and informal logic and the role they play in such an account of normativity. The proposal is that a normative claim cannot be expressed either in the terms of formal logical rules only or in terms of informal rules only. It is a combination of both. Hence, the way of determining an appropriate normative claim, namely, what kind of logical rules are normatively appropriate, on one side, and which doxastic state and deontic operator should be chosen, on the other, is a matter of mutual adjustment between requirements of rationality and human cognitive deductive reasoning set-up.

Springle, Alison (University of Pittsburgh)

On What Else Perceptual Representation Could Be

Abstract: Despite disagreements over various aspects of the nature of the content of perception, there's significant convergence among philosophers on I'll call the "Descriptive Conception" of perceptual representation. That so many philosophers have converged on the descriptive conception might seem to suggest that it is the only viable option for perceptual representation. On the other hand, it might suggest the presence of a blind spot; a poverty in the philosophical imagination. My aim in this short paper is to shed light on a hitherto unexplored conception of what else perceptual representation could be; to

broaden the philosophical imagination by arguing that what I call the "Instructive Conception" of perceptual representation is a viable naturalistic alternative to the descriptive conception.

Spychalska, Maria (University of Cologne), Viviana Haase (Ruhr University Bochum), Jarmo Kontinen (Ruhr University Bochum) and Markus Werning (Ruhr University Bochum)

Negation, prediction and truth-value judgments: evidence from ERPs.

Abstract: Prediction has been shown beneficial in various cognitive domains and argued to play a key role in linguistic comprehension [1,2]. Predicting sentence continuation as it unfolds seems straightforward in the case of affirmative sentences, but it is less clear how negation modulates prediction. First, it is generally more difficult to predict something that is absent. Second, the time-course of integrating negation into the semantic representation of a sentence is still not well-understood.

In our experiments we investigated how negation affects predictive processing. To this aim we compared the modulation of the N400 component in true and false, affirmative and negative sentences, such as "Julia has chosen/not chosen the plum". The N400 was time-locked to the presentation of content nouns (e.g. "plum"), whose cloze probability in a sentence was a function of the presented context scenario, i.e. of whether the scenario provided alternative referents that could be mentioned in the sentence. The scenarios contained three cards depicting three different objects, one or two of which were then selected (framed green) or rejected (framed red) by a virtual player. After the cards were marked, an affirmative or negative sentence (in German) was presented phrase-by-phrase: "X hat (nicht) den/die/das Y ausgewählt" ("X has (not) chosen Y"). In the first round of the experiments, all stimuli sentences provided a true description of the scenario (fillers of the opposite truth-value were added), while in the second round, all stimuli sentences gave a false description of the scenario. In the unique conditions, there was only one object that could complete a true (respectively false) sentence, i.e. one selected object in the affirmative condition or one rejected in the negative condition. In the multiple conditions, there were two alternative objects that could complete a true (respectively false) sentence.

It is well-established that the N400 is inversely correlated with the cloze probability of the triggering word. Our model assumes that the processor has to split the probability weights between the number of alternative objects available in the scenarios, leading to lower probability weights for objects that are presented with alternatives, and furthermore that the expected probability is higher for those objects that make the sentence true (truth-makers) rather than false (false-makers). In line with these assumptions, we observed larger N400 ERPs for words of lower probability weights (multiple conditions) than for words of higher probability weights (unique conditions), both for affirmative and negative sentences. This effect indicates that the processing of the unique referent is facilitated relative to the case where multiple referents are available in the scenario. Additionally, we observed a sustained positivity effect for the negative compared to affirmative conditions, which indicates possible reanalysis mechanisms related to the processing of negation. The effects were similar for true and false sentences: for true sentence the presence of multiple truth-makers increases the N400 amplitude, for false sentences it is the presence of multiple false-makers that increases the N400 amplitude.
[1] Kuperberg and Jaeger. Language, Cognition, and Neuroscience, 2015.
[2] Van Petten and Luka. International Journal of Psychophysiology, 2012.

Sreckovic, Sanja (Belgrade University)

Analyzing the results of experimental psychology: reasoning in non- and pre-linguistic creatures

Abstract: The presentation will focus on the conclusions derived from the experimental research on both human and animal cognition. The main thesis of the presentation will concern the relationship between the theoretical and the experimental work in this area. Certain behavior of some animal species is consistent with a capability for reasoning from an excluded alternative: the recognition that if there are only two possibilities and it is not one of them, it must be the other. This type of reasoning is one of the ways in which a creature can learn about what is not immediately perceptible. The possibility that such cognitive capabilities can exist without a capability for language is immensely important for many areas of cognitive science. Recently, there have been attempts to experimentally test whether some animals, as well as pre-linguistic human infants, are capable of this kind of reasoning. The experimental results typically do not hand us conclusive answers, so they have to be analyzed. I will examine certain assumptions behind the analysis of some of these experiments. I will focus mostly on the analysis of two studies. One study tested for the capability of reasoning by exclusion in great apes – Call's (2004) two cups task. Since the apes proved to be successful at the task, the results were taken by many commenters to indicate a capability for deductive reasoning. The apes supposedly reasoned: 'The food is either in A or B. It is not in A. Therefore it is in B'. However, in the subsequent theoretical research, there have emerged several other interpretations of the apes' behavior, all of which can explain the experimental results with fewer requirements placed upon their cognition. In order to distinguish between these interpretations, a more recent study was designed – Modey and Carey's (2016) four cups task was an extended, more demanding version of Call's task. The experiment was done with human children, and it was supposed to determine the exact mechanism underlying the exclusionary reasoning manifested in Call's experiment. It also proposed a way in which other species could be tested for the same cognitive capabilities. The children reasoned successfully at this task, and the researchers ultimately arrived at the deductive interpretation of their behavior. The aim of my presentation is to expose hidden assumptions behind their analysis of the experimental data, and to show that we cannot decide among the competing explanations based on the strategy they proposed. I also present possible
improvements for some of the interpretations, which should enable better predictive accuracy. This in turn exposes difficulties for further empirical testing, due to a lack of clarity concerning the differences between the competing accounts. The presentation concludes with a cautionary advice to experimental psychologists when attempting to use the empirical data to decide between the theoretically undeveloped theories.

Tan, Li Li (University of Cambridge)

Against the visual perception of high-level features

Abstract: There is an extensive debate on whether our visual experiences can represent more than just "low-level" features like shapes, colours, brightness and motion. Philosophers and psychologists have argued that various "high-level" features can also be part of visual phenomenology. These include natural kind features (e.g. having a tree appearance), social features (e.g. having a feminine appearance), and affordances (e.g. having a graspable appearance). This paper argues in favour of the view that high-level features are nothing over and above low-level features; they are identical to configurations of low-level features. This is because it is currently unclear how high-level features are supposed to make a difference

to visual phenomenology, and this should lead us to doubt that they must be introduced in order to provide a complete account of the phenomenology of human visual experience.

Taylor, Henry (University of Birmingham)

Two notions of kindhood: is disjunctivism inconsistent with psychology?

Abstract: This paper will consider the relevance of work in the philosophy of science on natural kinds to debates over disjunctivism. First, I introduce two different views of kindhood, which I label the schoolyard view and the jigsaw view. Second, I show that the schoolyard view is the dominant view of kindhood in the life and psychological sciences. Third, I show that if one accepts this view, then a commonly levelled criticism of perceptual disjunctivism can be answered. Specifically, on the view I recommend, though perceptual disjunctivism does turn out to be inconsistent with psychology, it is not in any way worrying for the disjunctivist or the psychologist.

Toribio, Josefa (University of Barcelona)

Accessibility, implicit bias, and epistemic justification

Abstract: It has recently been argued that beliefs formed on the basis of implicit biases pose a challenge for accessibilism, since implicit biases are consciously inaccessible, yet they seem to be relevant to epistemic justification. Recent empirical evidence suggests, however, that while we may typically lack conscious access to the source of implicit attitudes and their impact on our beliefs and behaviour, we do have access to their content. In this paper, I discuss the notion of accessibility required for this argument to work vis-à-vis these empirical results and offer two ways in which the accessibilist could meet the challenge posed by implicit biases. Ultimately both strategies fail, but the way in which they do, I conclude, reveals something general and important about our epistemic obligations and about the intuitions that inform the role of implicit biases in accessibilist justification.

Uusitalo, Susanne (University of Turku)

Vulnerability in addiction: a threat or a resource to agency?

Abstract: Vulnerability of addicted individuals is generally agreed upon: addicted individuals are a vulnerable population in health care, and in society in general. Their problematic use of psychoactive substances or certain types of problematic behaviours, e.g. in gambling, seem to take over their agencies. Partly for this reason, addicted individuals also struggle with battling against and overcoming their predicament. Indeed, it is common that people fail, repeatedly, in their attempts of recovery. After all, addiction is considered to be a disorder of impaired self-control and thus notoriously difficult to control by the individuals themselves.

Although vulnerability and addiction thus are closely linked—occasionally there even appears to be an implicit assumption that they are intrinsically related—there is different views on the source of vulnerability. For instance, the US National Institute of Drug Abuse (NIDA) states, "[b]rain sciences generate evidence of the vulnerability of addicted individuals referring to pathologies of the brain, typically induced by drug use", whereas the Global Commission on Drug Policy stresses that the vulnerability caused by drug policies is often more severe and harmful than the pathologies of the brain. In other words, both identify drug use as the loci of the vulnerability, but the cause for the vulnerability lies in different sources: The first

one attributes the cause to the effects of a psychoactive substance in the brain functioning, whereas the second attributes it to the illegal status of the drug and social policies, respectively. This difference can be understood as reflecting how addiction is discussed in frameworks that emphasize its biomedical nature over its social nature, and vice versa.

In the presentation I focus on how the agency of addicted individuals is seen as one involving vulnerability. Given its close link with addiction, vulnerability in addiction is typically taken to imply that addicted individuals' agency is undermined. This concurs with a common way of understanding vulnerability in general too. Nonetheless, instead of seeing vulnerability as something that can be identified with necessary and sufficient conditions, the two views on vulnerability mentioned above invite a more nuanced and multifaceted reading of the concept. In the presentation, I suggest that the vulnerability in addiction should be understood as relational; it seems to be context sensitive and in certain circumstances the same feature, e.g. drug use or gambling, that brought about vulnerability actually may enable and enhance the agency of an individual. Furthermore, I argue that the relational nature of vulnerability is particularly helpful in fleshing out the socially embedded and normative nuances of addicted individual's agency. Thus, even though vulnerability typically implies loss or lack of autonomy and/or agency, the relational reading of the concept of vulnerability sheds light on why using drugs as such or playing a game of chance as such do not make the agent inherently vulnerable in the sense that their agency is undermined. By doing so, the concept of vulnerability illustrates how the view of addicted individuals as "one cause" agents, driven by their object of addiction, is overly simplistic and implausible.

Vanello, Daniel (University of Geneva)

Mutual recognition, emotional engagement, and intentional communication

Abstract: A popular way of thinking about what renders human communication distinctive is via the theory of intentional communication. Intentional communication theory is committed to the claim that a necessary condition for genuine human communication to obtain is that the communicator must intend that its audience believe the conveyed information, and the audience's believing in the conveyed information must be at least in part due to the recognition that the communicator intended so. Intentional communication theory is also committed to the claim that the recognition of an intention essentially involves the inference of a mental state i.e. the intention, from overt and ostensive evidence. The commitment to such an inferential model of understanding other minds derives from the dominant idea in both philosophy and psychology that our ability to judge that another subject possesses mental states like oneself essentially involves inferring the presence of mental states from the other's overt ostensive behaviour. Within intentional communication theory, then, the mental states possessed by the communicator and the audience are specifiable independently of each other: intending that one's audience believe some bit of information can be specified independently from believing a bit of information in part because of recognising that the communicator intended so. And vice versa. Recently, some philosophers and psychologists have proposed an alternative way of conceiving the ability to understand that someone else possesses mental states like oneself. This alternative way is often branded the "second-person approach" because it relies on the notion of I-You relations of mutual recognition. The notion of mutual recognition denotes a mutual interdependence of the mental states involved when related to one another: I cannot think of you as a person with mental states unless at the same time you think of me as a person with mental states. The aim of this paper is to defend Vasu Reddy's mutual recognition model in terms of 'emotional engagement'. I believe that Reddy's account does not explain why the relation of mutual recognition should be necessarily emotional in kind and why such emotionality plays an essential role in

getting rid of the inferential model of understanding other minds in intentional communication theory. I want to fill in this gap. I argue that the mutual recognition emerges from the infant's signals to the mother which are motivated by the infant's desire to be responded to. The mother's subsequent response is both motivated by the desire to respond to the signals of the infant and are apprehended as such by the infant. Thus, the infant's recognition of the mother's response as motivated by wanting to respond to the infant's signals is both a genuine recognition of the mother's mental state and is a mental state of the infant that cannot obtain unless the mother genuinely is in that mental state, thus making the infant's mental state interdependent with the mother's. This form of mutual recognition is necessarily a form of emotional engagement because it is a relation of satisfaction or disappointment of one's desire to be responded to.

Veit, Walter (University of Bayreuth)

What can we learn from Models for the Evolution of Morality?

Abstract: Since Robert Axelrod's terminal work The Evolution of Cooperation, economists, biologists, philosophers and others have been eager to provide answers to question why we're not all selfish. Jason McKenzie Alexander tries to illuminate the requirements and potential threats to the emergence and spread of morality in his book The Structural Evolution of Morality. Analyzing a large scope of models and their robustness, he concludes that morality can be explained by combining evolutionary game theory, together with a theory of bounded rationality and research in psychology. Unlike Richard Joyce, Sharon Street and Alexander Rosenberg, he doesn't take this as an evolutionary debunking argument against morality. Instead he sees his work as providing an objective basis for morality in so far as moral behavior is in the long-term interest of everyone. This paper critically examines his work and concludes, that though his models provide a genuine advance upon work done by Axelrodd and Skyrms, his objective of arriving at objectivity morality must fail.

Venter, Elmarie (Ruhr University Bochum)

Modeling the Self

Abstract: In this paper, I focus on recent work by Jakob Hohwy and John Michael (2017) in which they provide an explanation of how to conceive of the self. They provide an account of the self much alike to that of Thomas Metzinger (2004, 2009) and claim that the self is identical to a set of endogenous causes modeled by the agent, this model is constantly updated just like any other generative model of the world. Claiming that the self model is updated in much the same way as a perceptual model of the world, prompts us to explore the role of the self in a perceptual state. Perceptual models of the world possess accuracy conditions involving the intentional content and the intentional mode that make up the intentional state. Recanati proposes that any proposition is meant to be evaluated with respect to the situation of the subject. He supports the distinction between self as subject and self as object. In this paper, I suggest that Hohwy and Michael's account neglects to explain the dimension of self as subject which is not updated as a model given the feature of entering the perceptual experience through the mode, and not through the content as in the case of self as object.

I argued that the agent as object of experience is firmly embedded in the world and capable of interacting with it in various sensorimotor and cognitive ways. In this sense, the self as object can enter the content of the perceptual experience in the same way as any other object of experience but still there is a lack of explaining how the self as subject enters this picture. In response to this, I proposed that the self as subject enters the conditions of satisfaction partly as a default parameter of the whole embodied organism. As

default parameter, the self as subject plays an implicit role in constructing generative models of the world and functions as an unarticulated constituent. This self as subject does not require identification as the self as object requires; it is therefore immune to the possibility of error through misidentification. The self as subject is captured through the process of active inference which is a state in which the organism is interacting with the world to minimize prediction error.

Vidmar, Iris (University of Rijeka)

Reacting to beauty: aesthetic consciousness, autonomy and agency

Abstract: Most discussions regarding the aesthetic judgments, or judgments of taste, focus on the question of whether such judgments are manifestations of subjective preferences, or whether they depict features that the objects objectively possess. Quite often, these discussions take a turn into a debate regarding the proper role of critics and critical judgments, and the question of whether one should modify one's own taste if one finds oneself disagreeing with others, most notably, the critics. In that sense, a theoretical framework within which the problem of aesthetic judgments is discussed consists of debates regarding realism vs. anti-realism of aesthetic properties, and regarding one's entitlement to one's aesthetic judgments and their subsequent justifiability.

As I argue in this paper, contemporary findings of cognitive aesthetics offer valuable and much needed solutions to both of these debates. To show that, I position the problem of aesthetic judgments against a slightly different framework, the one provided by the empirical approaches to aesthetic issues, most notably cognitive and evolutionary theories about development of human aesthetic preferences. More specifically, I show that traditional philosophical armchair theories of taste, particularly those advanced by Hume and Kant, were in fact anticipatory of some of the most contemporary conclusions advanced by cognitive aesthetics. I am primarily concerned with showing how some of Kant's most contentious claims get corroborated by contemporary findings of cognitive aesthetics, emphasizing four in particular: his notion of autonomous judgments, his idea about the empirical interest in the beauty and its connectedness to sociability, his notion of sensus communis, and his notion of the ideal of beauty.

At the core of my proposal is the notion of aesthetic agency, which I define as one's innate capacity to notice and react to the aesthetic features of the world and one's own experience of it, and to make judgments about these features that are grounded solely in one's being conscious of such experience, where it is rightly assumed that one is not mistaken about it. Theories developed by evolutionary scientists give convincing reasons to claim that such consciousness is hard-wired into our brains, and cognitive psychologists working on 'aesthetic' responses to art and environment suggest that aesthetic preferences are widely shared. This gives empirical force to Kant's view on sensus communis, but it raises two problems: that of explaining wide disagreements in aesthetic judgments, and the one concerning the sense that our aesthetic judgments are expressions of our most personal, inner, subjective selves, rather than our shared human nature. To solve these, I look at some of the factors that cognitive scientists underline as affecting our aesthetic preferences, and I use their insights to provide a more coherent account of aesthetic agency, which sees it as mediating between our personal, private selves, and our social interactions.

An addendum to my paper is a discussion on the methodological and theoretical value of introducing scientific research into philosophical debates, which, on some views, threaten to put philosophers out of business.

Werning, Markus (Ruhr University Bochum) and Erica Cosentino (Ruhr University of Bochum)

The Interaction of Bayesian Pragmatics and Lexical Semantics in Sentence Meaning Composition

Abstract: It has been widely acknowledged that a preceding discourse can influence the way sentence meaning is composed from lexical meaning. A prominent view is that the contextual influence is mainly due to the semantic similarity between parts of the discourse context and the words in the target sentence. It is however highly controversial whether also pragmatic aspects of the discourse context other than the mere resolution of indexicals and anaphors can immediately affect sentence meaning composition.

To investigate the contrast between a Semantic Similarity and a Free Pragmatic account we will look into the way subjects make probabilistic predictions on the completion of a sentence given a preceding discourse. Quantitatively, the semantic similarity can be determined by Latent Semantic Analysis (Landauer & Dumais, 1997), whereas we will use the framework of Bayesian Pragmatics (Frank & Goodman, 2012) to calculate the pragmatic influence – in particular, concerning the rationality of the speaker's intentions in a narrative. As a model of lexical structure we apply the Generative Lexicon approach (Pustejovsky, 1995).

Wilson, Keith (University of Edinburgh)

Windows on Time: The Temporal Structure of Multisensory Experience

Abstract: Perceptual experience, unlike remembering or imagining, is characteristically an experience of how things are in the present moment, or 'now'. However, each of our sensory modalities operates on different timescales, with stimuli taking differing lengths of time to reach the relevant receptors, and incurring distinct processing delays. In order to create a unified experience of the present, the brain must bind together stimuli processed in different sensory streams as occurring 'at the same time' — the so-called temporal binding problem (Pöppel & Bao 2014). One mechanism that has been proposed as to how this is done claims that perceptual processing is organised into a series of 'temporal windows', each having a minimal (e.g. 30–60ms) duration out of which longer experiences are composed. The existence of such minimal units, however, is controversial and in tension with certain philosophical views of temporal experience which either deny that such units exist (Dainton 2000), or that they have explanatory precedence (Soteriou 2013, Phillips 2014). In this paper, I evaluate the empirical evidence for the existence of temporal windows, and their upshot for the philosophy of temporal experience.

The first part of the paper sets out the evidence in favour of the temporal windows hypothesis, the most striking of which is Pöppel's (1970) claim of a 30ms periodicity in reaction times in an audio-visual task — a finding disputed by Ulrich (1987), but subsequently replicated by Dehaene (1993). Various possible explanations for this effect are considered, including that the resulting periodicity may be a result of motor action rather than perceptual processing. Other studies find a lower threshold for simultaneity judgements, with the precise threshold varying according to the sensory modality, or modalities, in question. This suggests that if there is a periodicity effect, it is modality-specific, or that the events within the resulting 'functional moments' are experienced as possessing some form of internal temporal ordering and not as being simultaneous. The evidence for this, however, is inconclusive, suggesting a need for further empirical research.

The second part of the paper examines evidence from the multisensory processing literature which

suggests the existence of a longer (approx. 250ms) window across which we are able to bind temporally discrepant cross-modal stimuli (Harrar, Harris & Spence 2016). Unlike Pöppel's functional moments, multisensory binding preserves internal temporal order and is relatively consistent across modalities, suggesting the existence of a higher-level mechanism. The differing features and durations of the resulting temporal windows suggests that not only are they distinct, but that they may form part of a hierarchy whose 'levels' have distinct functions and temporal characteristics. This lends empirical support to William James' notion of the "specious present" (1890) and suggests that, at the fundamental processing level, perceptual experience may consist in a series of discrete 'moments' rather than being strictly continuous. This presents a prima facie challenge to philosophical views that deny the existence of any such minimal units of experience, or that they have explanatorily precedence over longer temporal slices.

References

Arstila, Valtteri, & Dan Lloyd, eds. (2014). *Subjective Time: The Philosophy, Psychology, and Neuroscience of Temporality*. Cambridge, MA: MIT Press.

Dainton, Barry (2000). *Stream of Consciousness: Unity and Continuity in Conscious Experience*. London: Routledge.

Dehaene, Stanislas (1993). 'Temporal Oscillations in Human Perception'. *Psychological Science* **4** (4): 264–70.

Harrar, Vanessa, Laurence Harris & Charles Spence (2016). 'Multisensory Integration is Independent of Perceived Simultaneity'. *Experimental Brain Research* **235** (3): 763–75. DOI: 10.1007/s00221-016-4822-2.

James, William (1890). *The Principles of Psychology*. New York: Dover.

Phillips, Ian (2014). 'The Temporal Structure of Experience'. In V. Arstila & D. Lloyd, eds, pp. 139–58.

Pöppel, Ernst (1970). 'Excitability Cycles in Central Intermittency'. *Psychol. Forsch.* **34**: 1–9.

Pöppel, Ernst, & Yan Bao (2014). 'Temporal Windows as a Bridge from Objective to Subjective Time'. In V. Arstila & D. Lloyd, eds., pp. 239–306.

Soteriou, Matthew (2013). *The Mind's Construction*. Oxford: Oxford University Press.

Ulrich, Rolf (1987). 'Threshold Models of Temporal-Order Judgments Evaluated By a Ternary Response Task'. *Perception & Psychophysics* **42** (3): 224–239.

Wolf, Julia (Ruhr University Bochum)

Modelling the Paradox of the False Belief Task Using Mental Files

Abstract: It is a well replicated finding that children as young as 15 months are able to pass the implicit false belief tasks (Scott & Baillargeon, 2017). These findings seem to indicate some early understanding of other people's beliefs. Nonetheless, linguistically quite well developed children still fail the explicit false belief task till they are approximately 4 years old (Wellman, Cross, & Watson, 2001). This is known as the paradox of the false belief task (De Bruin & Newen, 2012). There are two main aims of this paper, firstly, we argue that evidence from the active helping paradigm of the false belief task should be interpreted as suggesting that there is an intermediate stage in false belief understanding between implicit and explicit false belief understanding. Secondly, we suggest a new account of this paradox in terms of an interplay

between both cognitive and situational factors. Making use of the Mental Files framework, recently put forward by Perner and Leahy (2016), we develop a more detailed account of cognitive development, which is also able to provide an account of findings from the active helping behaviour paradigm (Buttelmann, Carpenter, & Tomasello, 2009). We argue that we also need to consider the role of situational factors and how these relate to cognitive development in order to fully explain the paradox of the false belief task. In particular, we argue for the importance of situational factors which highlight the perspective of another person. Moreover, we argue that situational factors play an important role in the development of cognitive abilities themselves, namely by providing the origin of perspectival thinking and thereby generating the need to relate different perspectives to reality for action.

Yamamoto, Naoko (Nara University)

An Experimental Analysis of Children's Misunderstanding of Tautology

Abstract: Many researchers have experimentally examined children's (mis)understanding of figurative expressions (e.g. metaphor and irony) and argued that irony is more difficult for children to understand than metaphor. No one has yet discussed children's (mis)understanding of tautology *A is A* (≒ Japanese *A wa A* (*da*)). This paper claims, based on experimental data on Japanese children, that characteristics of children's misunderstanding of tautology are common to ones observed in metaphor and irony, and proposes that the interpretation of tautology is a process that requires higher processing competence.

This study involves 193 first-sixth year students in elementary school (aged 6-12) and 20 university students. Participants were instructed to carefully read a booklet of four tautologies, which were each presented with contextual information and three options as its interpretation(s), and then judge whether each option is appropriate for an interpretation of the expression. The first-sixth grades have 32%, 39%, 35%, 42%, 51%, 44% average accuracy percentages of the tautologies, respectively, which are at lower rates than that of the adults', 85%. This shows that schoolchildren cannot understand tautologies properly.

(1) [A six-year-old girl mutters about her younger sister.]

Imoto wa itsumadetattemo imoto. (word-for-word translation: Little sister is always little sister.)

a. My little sister won't change. / b. A little sister is a younger sister who has the same parents. / c. The six-year-old-girl treats her sister like a baby.

(2) Son (Takashi): When will you play with me?

Father: *Ato wa ato da.* (word-for-word translation: Future is future.) Play with your brother.

a. Takashi will play with his father after enjoying himself with his brother. / b. His father does not want to say when he can play with Takashi. / c. His brother will play video games with Takashi.

The girl's muttering in (1) communicates her sister's unchangeability such as (1a). The children show a strong tendency to choose (1b), which is a lexical meaning of *imoto* (≒ English *little sister*). They cannot extract relevant information like (1a) from associated encyclopedic information about *imoto* depending on the context. This failure is similar to one found in children's understanding metaphor and irony (cf. Winner 1988). This shows that <u>when interpreting a tautology, children prioritize a literal meaning like metaphorical and ironical cases.</u> In (2), the father communicates his intention of refusing his son's request to play with

him, as shown in (2b). The children are prone to choose (2a). The answer seems to be derived from encyclopedic information about *ato* (≒ English *future*), but it is not suitable for the context. The children fail to read the speaker's intention such as rejection. This error is common to one in children's irony understanding (cf. Matsui 2013). This shows that <u>when comprehending a tautology, children find it hard to read a speaker's intention communicated by the expression like ironical cases</u>. These findings reveal that the characteristics of children's misunderstanding of tautology are common to ones found in metaphor and irony. It seems therefore that the process of interpreting tautology requires higher processing competence.

Selected References

Matsui (2013) *Kodomo no Uso, Otona no Hiniku* (*Children's Lies, Adults' Ironies*), Iwanami Publishing.

Winner (1988) *The Point of Words*, Harvard.

Young, Nick (University of Milan)

Against Hearing Sounds

Abstract: Here I argue that we do not hear sounds, but rather sound sources: events in which material objects move or interact in such a way as to vibrate and produce sound waves in the air. On one type of view sounds are individuals produced by sound sources, on another they are properties of source events or of the objects participating in source events. Proponents of both types of view agree that we hear sound sources through hearing sounds. I argue that this type of mediate perception faces difficulties, and that a better account of source hearing can be arrived at if give up the idea that audition must involve hearing sounds. I go on to give details of how soundless audition should be understood, and then respond to two types of argument that could be made in defence of sounds.

Zilhao, Antonio (University of Lisbon)

Cognition and Evolution: The Case of Means-End Reasoning

Abstract: In "The Evolution of Means-end Reasoning", Papineau claims that means-end reasoning played a key role in the evolution of human cognition. His substantive characterization of means-end reasoning proceeds along the lines of a common philosophical reading of the classical rational choice framework. Papineau also puts forth the following sub-claims: 1. Means-end reasoning is a peculiarly human ability; 2. Means-end reasoning is a biological adaptation (rather than a spandrel); 3. Means-end reasoning is a modular mechanism added on by evolution to other fast and frugal mental mechanisms.
Two poles (A and B) mark out the theoretical space within which the debate around the nature of action-oriented mental mechanisms takes place. Pole A theorists claim that human cognition is mainly sub-served by heuristics; they deny the legitimacy of characterizing human action-oriented mental mechanisms in terms of means-end reasoning. Pole B theorists share with Papineau a folk-psychological/rational choice approach to cognition and decision. But they claim humans to be endowed with a general-purpose means-end reasoning mechanism. They disagree with sub-claim 3.
Following in the wake of Papineau's suggestion, it is tempting to look for a solution to this debate by introducing the diachronic dimension and going evolutionary. A number of evolutionary arguments have been advanced in this context.
The evolutionary justifications supporting pole A views are the following: 1. The Trade-off Argument: Sacrificing accuracy for speed and/or frugality is a trade-off promoting the fitness of evolved organisms,

humans included. 2. The Modularity Thesis: Cognition is modular (because modularity facilitates speed and frugality); 3. The Heuristics Thesis: fitting in rather well within a modular architecture, heuristics are the building blocks of cognition.

The evolutionary arguments supporting pole B views rely on the following premises: 1.It is possible to imagine selection environments rewarding accuracy rather than speed and frugality. 2.Human social environment is one such environment. 3. Social environment played a crucial role in the evolution of humans. 4. Obtaining accuracy is best served by the classical decision-theoretic approach.

What separates these arguments is the fact that they base their conclusions on the appeal to a different kind of evolutionary environment. But we know next to nothing about the environments our ancestors actually evolved in. Thus, we are in no position to adjudicate between them. After deciding to go evolutionary, we reached a stalemate.

However, theorists belonging to the Ecological Rationality Group (ERG) claim having found a disclaimer against premise 4. They report having performed experiments with humans having shown that distinct fast and frugal heuristics fare better than rival classical approaches in the accuracy dimension. Assuming they are right, both the classical view and Papineau's middle ground position lose out. But are they right? Nobody knows.

After having gone through the evolutionary arguments, we realized that what we are left with is a synchronic debate in Cognitive Psychology. So far, the evidence collected in this debate is inconclusive. However, this synchronic inconclusiveness extends itself diachronically. Thus, contrary to Papineau's expectations, the introduction of the evolutionary dimension in this debate does not tip the scales.

Accepted symposia

Cognitive Penetration: A fresh look at the debate with new experimental results

Albert Newen (Ruhr-Universität Bochum), Christoph Teufel (Cardiff University), Valtteri Arstila (University of Turku), Petra Vetter (University of London)

To what extent is our perceptual experience influenced by our concepts, beliefs, desires? Cognitive penetrability describes the influence of high-level cognitive factors on perceptual experience and has become a hotly debated topic in philosophy of cognition. What is the phenomenon in question: cognitive penetration (in a narrow sense) describes an influence of the content of higher cognitive states (e.g. the activation of a concept, belief etc.) on the percept that is caused by a given sensory input while the state of the sensory organs (in terms of spatial attention and sensory input) is held constant. What is held constant is: 1. the object or scenario causing the visual input, 2. the perceptual conditions, 3. normally functioning sensory organs and 4. the absence of attentional shifts. If an influence of a cognitive content (e.g. activation of a concept or belief) can nevertheless change the perceptual experience, this is a case of cognitive penetration. If cognitive penetration exists, we have to give up a strict demarcation between perception and cognition which seems to be necessary to presuppose to explain stable visual illusions. How can we make progress in clarifying whether cognitive penetration exists or not and which theoretical consequences this has? This symposium aims to develop a new look on the debate by presenting new experiments which are relevant for the phenomenon in question and by (a) having a look on the most challenging experiments so far discussed, (b) by especially presenting new own relevant experiments which allow us to widen the perspective and (c) ideally by outlining new theoretical perspectives.

Communication, commitment, and epistemic vigilance

Mazzarella, Diana (University of Neuchâtel), Christophe Heintz (Central European University), Thom Scott-Phillips (Central European University), Francesca Bonalumi (Central European University), Margherita Isella (Università Vita-Salute San Raffaele) and Julius Tacha (University of Vienna and Central European University).

Communication is an essential source of knowledge about the world: much of what we know we have learned from others via testimony. Communication, however, is open to the risk of accidental or intentional misinformation and, consequently, the acquisition of beliefs via testimony requires careful calibration of trust and vigilance (Sperber et al., 2010).

The aim of this symposium is to address the question of the relation between trust and vigilance by bringing together contributions from different fields of research: evolutionary psychology, linguistics, and experimental social psychology.

The symposium will devote special attention to the mechanism of 'commitment' in communication.

Commitment allows communicators to gain the trust of the audience and increases the chance that their message will be accepted as true. This comes at a cost: committed communicators incur additional reputational costs when the message is found to be unreliable (Vullioud, Clément, Scott-Phillips & Mercier, 2017). That is, commitment puts at risk reputation and future opportunities for cooperation.

The first contribution, [1], investigates the relationship between reputation and commitment and it examines its workings within the domain of communication. The claim is that the ostensive – intentional – nature of human communication brings about a specific kind of commitment: the communicator's commitment to the relevance of her communicative act.

The second contribution, [2], discusses the co-evolution of communicative abilities and epistemic vigilance capacities. The aim is to consider some of its implications for theories of human communication and to provide a novel account of its zoological distinctiveness and expressivity.

Finally, the symposium addresses the question of how commitment is modulated in communication by looking at both linguistic factors [3] and interpersonal factors [4]. The third contribution, [3], provides empirical evidence that commitment is modulated as a function of the manner in which information is conveyed and belief is propagated. Commitment is weaker when the message is implicated rather than asserted or presupposed. The last contribution, [4], presents a series of studies examining the role of 'reliance' in commitment modulation. The more the audience is expected to rely on the communicated message, the stronger the commitment attributed to the communicator.

The symposium combines theoretical and experimental approaches and it crosscuts disciplinary boundaries to investigate the complex interplay between trust and vigilance in human communication. Well-adjusted trust must take into account the extent to which communicators commit to their messages in a fine-grained way. This symposium aims at shedding new light on how this is made possible.

Epistemic Innocence

Puddifoot, Katherine (University of Birmingham), Sophie Stammers (Kings College, London), Ema Sullivan-Bissett (University of Birmingham), Valeria Motta (University of Birmingham) and Matilde Aliffi (University of Birmingham)

Abstract: The notion of epistemic innocence was introduced to capture how cognitions that are irrational or inaccurate can nonetheless bring significant benefits which could not have otherwise been attained (Bortolotti 2015, 2016; Bortolotti and Miyazono 2015; Bortolotti and Sullivan-Bissett forthcoming; Letheby 2016; Puddifoot 2017; Sullivan-Bissett 2015, forthcoming). The benefits that the notion is used to pick out are epistemic ones, associated with the ability to gain true belief, knowledge and understanding. A cognition is epistemically innocent if it is epistemically faulty in some way, for example, irrational or inaccurate, and operates in ways that could increase the chance of knowledge or understanding being acquired, where alternative, less costly cognitions that bring the same benefits are unavailable. Recently, the application of the notion of epistemic innocence has been expanded to include cognitive mechanisms (Puddifoot and Bortolotti 2018). This symposium showcases various ways that the innovative notion of epistemic innocence can be applied to cognitions and cognitive mechanisms and shows how it can also be applied to emotions, such as those involved in loneliness. The contributors thus show innocence in action, and the importance of such a notion in our projects of epistemic evaluation.
Sophie Stammers argues that confabulations are produced by a narrative capacity that is epistemically

innocent. Valeria Motta and Matilde Aliffi focus on seemingly irrational loneliness, which does not reflect the actual quantity of a person's social interactions, arguing that there are epistemic benefits to this loneliness, so it can be epistemically innocent. Ema Sullivan-Bissett focuses on monothematic delusions, arguing that opposing empiricist theories of delusion formation can converge on the epistemic innocence of these cognitions, and thus have the notion at their disposal. Katherine Puddifoot argues that the notion of epistemic innocence can be used to establish that eyewitnesses who are uninformed by psychological findings are likely to give less credence than they should to some eyewitness testimony.

How the science of action can illuminate the philosophy of action (and beyond)

Brozzo, Chiara (University of Tübingen & Centre for Integrative Neuroscience), Hong Yu Wong (University of Tübingen & Centre for Integrative Neuroscience), Gregor Hochstetter (University of Tübingen & Centre for Integrative Neuroscienceand), Krisztina Orbán (University of Tuebingen)

Abstract: Can the science of action illuminate the philosophy of action, and, if so, how? It may be thought that the two do not (or, even, should not) interact at all. In response to this question, we'll start with a theoretical defense of the possibility for the science of action to contribute to the philosophy of action. This defense will counter arguments based on the difference between the space of reasons and the space of causes. We will then offer a number of examples of a fruitful interaction between the philosophy of action and the science of action.
Philosophy of action will be interpreted in the broad sense, meaning not just answers to the question as to what distinguishes actions from mere happenings, but also other areas of philosophy that study action. One example is social cognition, which involves the study of how observed actions are understood.
We will start with a challenge to standard accounts of action that derives from the science of action: how do we distinguish everyday actions from bodily movements that are not actions, such as those produced by Anarchic Hand Syndrome patients? We will continue with the question as to whether understanding others' actions involves mindreading, i.e. the ascription of propositional attitudes. Here data from the science of action will motivate investigating an interesting middle ground between the following two extremes: that action understanding does not involve mindreading at all and that action understanding involves the ascription of standard propositional attitudes. We'll end on an example of a surprising and unprecedented example of how the science of action can illuminate even areas of philosophical investigation that do not concern action at all. In particular, it will be shown how the notion of first person reference relies on the key ethological and developmental phenomenon of self-directed action.

Levels of classification and explanation in clinical psychology and psychiatry

Lalumera Elisabetta (University of Milano-Bicocca), Maria Cristina Amoretti (University of Genoa), Luca Malatesti (University of Rijeka), Marko Jurjako (University of Rijeka), Inti Brazil (Radboud University), Markus Eronen (University of Groningen) and Annemarie Kalis (Utrecht University)

Abstract: The aim of the symposium is to discuss a recent important proposal on the nature of mental disorders, elaborated by Denny Borsboom, a psychologist and psychometrician, Angélique Cramer, a clinical psychologist, and Annemarie Kalis, a philosopher, in a series of articles appeared or forthcoming in journals such as *World Psychiatry* and *Behavioral and Brain Sciences*. Their positive idea is that mental disorders are networks of interrelated symptoms, and the negative underpinning is that biological and neurological reductionism should be avoided in psychiatric research. Along with a presentation of the target proposal, the symposium will comprehend three comments. One of these comments is by two psychologists, one by

two philosophers of psychiatry and psychology, one by two philosophers of psychology and psychiatry and a cognitive neuroscientist.

Making things up: Pretending, Fabricating, and Confabulating. Interdisciplinary and Contextual Analysis

Rucińska Zuzanna (Polish Academy of Sciences and Leiden University), Krystyna Bielecka (University of Warsaw), Monika Chylińska (John Paul II Catholic University of Lublin), Arkadiusz Gut (John Paul II Catholic University of Lublin), Adrianna Smurzyńska (Jagiellonian University) and Piotr Litwin (University of Warsaw)

Abstract: Pretending, fabricating and confabulating are all instances of 'making things up'. Pretending is characterized as behaving so as to make it appear that something is the case when in fact it is not; it is making things up in fiction, often for the purposes of play. Fabricating is characterized as inventing or constructing facts or stories often for the purposes of deception (including self- deception). Confabulating is characterized also as making things up; the term functions mainly in the clinical context where it relates to fabricating stories and imaginary experiences as a compensation for loss of memory.

This symposium will focus on elucidating the various environments in which these forms of 'making things up' occur, in order to bring out their similarities and differences in a philosophical analysis and in a variety of psychological settings: developmental, experimental and pathological ones. It will include an analysis of different contexts for pretending, fabricating and confabulating, from making things up in everyday playful interactions, where fabricating is partly conscious and controlled, to unconscious fabricating in the response bias found in experimental settings, or uncontrolled confabulating in Korsakoff's syndrome. The questions of what is being made up, why, and how, will be targeted in more depth, including a discussion of the role of scripts, representations, simulation, mentalization, embodiment and narrative practices in the explanation of these phenomena. Our individual topics include an analysis of pretend engagements with a focus on the role of narrative practices (topic 1), fabricating and recreating scripts and event schemas in children's pretend play (topic 2), comparison of everyday confabulating to clinical confabulating (as seen in Korsakoff's syndrome) in the process of self-representation (topic 3), simulating and fabricating in the process of understanding other minds (topic 4) and the role of fabricating in the unconscious response bias found in some experimental settings (topic 5).

This symposium thereby brings together a variety of philosophical perspectives and psychological approaches to understanding the conceptual background of making things up, in the context of developmental, experimental and clinical findings. We aim to show in practice how philosophical theories and psychological studies can be brought together to enrich our understanding of the complex nature of pretending, fabricating, and confabulating.

Time and Choice

Christoph Hoerl (University of Warwick), Ruth Lee (Queen's University Belfast), Teresa McCormack (Queen's University Belfast), Patrick Burns (Queen's University Belfast), Patrick O'Connor (Nottingham Trent University), Alison Fernandes (Trinity College Dublin), Fabio Paglieri (ISTC-CNR Roma), Marco Marini (London School of Economics), Dieter Thoma (University of St Galen), Agnieszka Ewa Tytus (University of Mannheim)

Abstract: Time has an influence on our decisions and preferences that has attracted the attention of both philosophers and psychologists. Two aspects of this influence have been of particular interest across the disciplines. The first is humans' propensity to discount the value they place on events in the non-immediate vs. the immediate future: people often opt for smaller sooner rewards instead of waiting for larger later ones. The second is people's tendency to differentially value certain events depending on whether they are located in the past versus the future. Philosophers have been particularly interested in the rationality of both these phenomena and the implications of the phenomena for ethics and accounts of personal identity. There is also philosophical debate over how asymmetries of attitudes towards the past and future may connect with metaphysical claims about the nature of time itself. Separately, psychologists and behavioral scientists have extensively explored the cognitive processes underpinning these phenomena and the variables by which they may be manipulated. Moreover, recent research in linguistics has examined how the grammatical characteristics of language may impact on intertemporal choice. This symposium looks at questions about intertemporal choice from the perspective of philosophy, experimental psychology, linguistics and developmental psychology.

Hoerl looks at intertemporal choice in the context of broader questions about the nature of temporal reasoning. In a wide range of decision-making tasks, evidence of signature processing limits and performance contrasts has supported the notion that many cognitive processes may be a function of two interacting systems, one that is heuristic, associative and immediate, and another that operates more slowly and analytically. Hoerl argues that each may contribute uniquely to the mitigation of temporal discounting.

Marini and Paglieri investigate the roles of time pressure and immediacy on the influence of decoys in intertemporal choice, suggesting that a comparative process is required in order for decoys to bias decisions and that such influence may be possible in the sooner-and-smaller (as opposed to later-and-larger) reward direction only when rewards are not presented as immediate.

Thoma and Tytus investigate recent claims that individuals' intertemporal choices are influenced by how strongly the difference between present and future is marked grammatically in their first language. They present evidence from a cross-linguistic study suggesting that, while natively speaking a language entailing overt grammatical dissociation of the past from the future may be associated with a psychological distance that influences intertemporal choice, the direction of this association may depend on the ecological validity of the tasks presented to participants.

Lee, McCormack, Burns, O'Connor, Fernandes, and Hoerl investigate the extent and nature of a bias toward valuing the future in early and middle childhood, presenting preliminary evidence that an adult-like bias toward preferring pleasure to be in the future rather the past, and pain to be in the past rather than the future, begins to emerge between the ages of 5 and 7.

Tool innovation, manufacture, and use in evolutionary, ontogenetic, and neurocognitive perspectives

Hohenberger Annette (Middle East Technical University), Gökhan Gönül (Middle East Technical University), Alice Auersperg (University of Vienna), Sarah Beck (University of Birmingham), Clare Whalley (University of Birmingham), Nicola Cutting (University of Birmingham), Louise Bunce (Oxford Brooks University), François Osiurak (University of Lyon), Emanuelle Reynaud (University of Lyon), Jordan Navarro (University of Lyon) and Mathieu Lesourd (Aix Marseille University)

Abstract: Humans use tools as interface devices, extending the range of their action possibilities and increasing the degree of efficiency of manipulating their environment. The multifarious functions which tools can serve, goals they can help to reach, and potentiality of application have made tools indispensable, beneficial devices for their users. Being situated at the nexus between body anatomy, object physics, cognition, language, sociality, and culture, tool-related behaviors are an exquisite topic amenable to interdisciplinary studies. Tool-innovation, manufacture, and use has attracted renewed attention from a variety of disciplines and perspectives, recently. Comparative studies have shown that not only adult humans but also many animal species show tool-related behaviors in captivity as well as in the wild. From developmental studies we know that children are capable of using tools from early onwards, however, that they experience difficulties innovating tools. Neurophysiological studies reveal that different aspects of tool use are supported by different brain areas, comprising temporal, parietal, and frontal lobes. Yet, many aspects of tool-related behaviors are still poorly understood. In this symposium we address the topic from evolutionary, ontogenetic, and neurocognitive perspectives – covering three different time-scales: macro (evolution), meso (ontogeny), and micro (brain processes).

The first paper by Auerspach presents novel research on captive cockatoos' ability to spontaneously manufacture tools – although these birds are not specialized for complex tool-related behaviors. Animals were observed to break sticks from larch wood, cutting cardboard stripes with their beak, defoliating twigs or (un-)bending a wire to rake distal food items into reach. These findings are insightful as they shed light on the evolutionary origins of tool-related cognition.

The second paper by Beck, Whalley, Cutting, & Bunce studies the effect of social and psychological context on 4-5 and 6-7-year-old children's tool-innovation abilities, using the "hook task". In Study 1, the authors show that social context, namely being left alone, facilitates older children's success at tool-innovation. In Study 2, they show that psychological context, namely priming children with distant objects, facilitates more abstract thinking, thus promoting their tool-innovation success.

The third paper by Gönül and Hohenberger comprises a set of four studies. The authors show that social demonstration of tool making, familiarity of the tool, collaboration among children in dyads, as well as cultural aspects may facilitate 3-6 and 5-7-year-old children's tool-making success. Moreover, hierarchical structuring can predict tool-making in particular conditions.

The fourth paper by Osiurak, Reynaud, Navarro, & Lesourd introduces the "Technical Reasoning Hypothesis", arguing that reasoning about physical object properties is essential for tool-use. They adduce evidence from patient and neuroimaging studies implying a crucial role of the left inferior parietal cortex. Most interestingly, their latest research indicates that observing others using tools also instigates active reasoning processes about the observed mechanical actions.

In the subsequent discussion with the audience we will highlight possible points of convergence among the various perspectives taken in this symposium and seek to identify promising future lines of interdisciplinary research on tool-related behaviors.