The text that follows is an author's preprint of a paper published in *Journal of Experimental and Theoretical Artificial Intelligence*, December, 2018, DOI: 10.1080/0952813X.2018.1544285

A mosaic of Chu spaces and Channel Theory II: Applications to Object Identification and Mereological Complexity

Chris Fields
23 Rue des Lavandières
11160 Caunes Minervois, FRANCE
fieldsres@gmail.com

and

James F. Glazebrook
Department of Mathematics and Computer Science
Eastern Illinois University, 600 Lincoln Ave.
Charleston, IL 61920–3099, USA
jfglazebrook@eiu.edu
Adjunct Faculty
Department of Mathematics
University of Illinois at Urbana–Champaign
Urbana, IL 61801, USA

November 28, 2018

Abstract

In this Part II of a two-part work, we proceed from the survey of concepts and techniques of Chu spaces and Channel Theory in Part I to the characterization of human visual object identification, beginning with the construction of uncategorized object files and proceeding through categorization, individual object identification and the tracking of object identity through time. We investigate the relationship between abstraction and mereological categorization, particularly as these affect object identity tracking. This we accomplish in terms of information flow that is semantically structured in terms of local logics, providing an inferential mechanism for object identification and tracking. We introduce categorical Cone-Cocone Diagrams to explicitly capture a bidirectional duality between "token" roles and "type" roles, and show that all representations can be considered to play both roles. We discuss the emergence of mereotopology from the representation of classifications with underlying simplicial complexes, and briefly explore the emergence of geometric relations and interactions between objects. Throughout we discuss the empirical support for the utility of this descriptive mechanism, particularly as a scale free organization that is applicable to cognition and to AI systems in general, and raise open issues and problems.

Keywords: Chu space, Channel Theory, Categories, Cognition, Perception, Object-Event

Files, Object Token, Context, Episodic Memory, Cone-Cocone Diagram, Mereological Complexity.

Contents

1 Introduction

In Fields and Glazebrook (2018) we surveyed the categorical methods for describing semantic information flow provided by Chu spaces (Barr, 1979, 1991; Pratt, 1995, 1999a,b) and Channel Theory (Barwise and Seligman, 1997; Seligman, 2009). In this second part of a two-part work, we apply these concepts and tools to formulate and explore a very general hypothesis: that the neurocognitive architecture employed by humans – and possibly all mammals – is both structurally and functionally scale-free. Evidence that this hypothesis may be true, at least as a good approximation, comes primarily from two sources: 1) graph-theoretic analyses of functional neuroimaging data that indicate small-world or "rich-club" structure at multiple scales (Bassett and Bullmore, 2006; Sporns and Honey, 2006; Glazebrook and Wallace, 2009; Rubinov and Sporns, 2010; Sporns, 2013), and 2) the theoretical coherence and explanatory power of multi-layer, expectation-driven, recurrent processing methodologies, particularly the adaptive resonance (Grossberg, 1980, 1988, 2007, 2013), and free-energy/predictive-coding (Friston, Kilner and Harrison, 2006; Spratling, 2008; Friston and Kiebel, 2009; Friston, 2010; Bastos et al., 2012; Spratling, 2016, 2017) frameworks. Categorical methods provide an attractive avenue for exploring the hypothesis of a scale-free cognitive architecture in general terms, as they are both expressively richer than graph theoretic methods and more general than any specific recurrent-processing model. The particular tools and concepts provided by Chu spaces and Channel Theory are, moreover, designed for this application. While these tools have been applied to human, as opposed just to abstract, cognition as reviewed in Fields and Glazebrook (2018), these applications have been neither general nor strongly driven by empirical results. Our objective here is to take this previous work to the next level, showing that the formalism of Chu spaces and Channel Theory provides useful insights into the empirically-characterized, systems-level architecture of human cognition.*

As a model system for human cognition more generally, we focus on the visual identification of individual objects as time-persistent entities. Visual object identification naturally involves multiple scales, from local analysis of image components such as edges or gradients to the semantic interpretation of sequences of scenes as revealing causal processes or goal-oriented actions. Visual object identification is one of the best-studied of all cognitive processes, although deep issues concerning long-term persistence judgments remain unresolved (for recent reviews see Fields, 2016, 2017). It involves sensory, attention, memory, affective, and motor systems working in coordination, and hence touches on most major processing pathways involved in general cognition. We focus in particular on developing category-theoretic descriptions of three interdependent components of visual object identification: 1) the construction of object files (Kahneman, Triesman and Gibbs, 1992) and object tokens (Zimmer and Ecker, 2010), 2) the binding of type and token information in object categorization (Martin, 2007; Keifer and Pulvermüller, 2012), and 3) the recognition and categorization of mereologically-complex individuals. We characterize each of these component processes in terms of both bottom-up and top-down information flows semantically structured by the 'object-attribute' relations of Chu spaces and/or the 'token-type' Classification relations of Channel Theory. In doing so, we combine these methods with the simplicial-topological and categorical techniques that were developed in Fields and Glazebrook (2018). We show that bottom-up and top-down processes can be considered category-theoretic duals, and that this duality structure is preserved across multiple scales of both type abstraction and mereological complexity.

While the notion of "entry-level" categories and the extension of such categories both upward and downward in an abstraction (or type) hierarchy has been intensively investigated both experimentally (Clarke and Tyler, 2015) and theoretically (Sowa, 2006), the representation of mereological (i.e. part-whole) complexity has received far less attention. Mereological categorization can be functionally dissociated from abstraction-based categorization in humans, e.g. in high-functioning autism where "weak central coherence", and hence deficit understanding of mereological complexity may be displayed alongside normal or even superior abstraction ability (Happé and Frith, 2006; Booth and Happé, 2018). How abstraction-based types and mereological types are related, and how their implementations in humans are related, thus remain to be worked out. We focus on one component of this issue: the question of how tokens representing individual, re-identifiable objects, i.e. object tokens as defined in Zimmer and Ecker (2010), are able to participate both as such, and as instances of classified types in both hierarchies simultaneously. How, for example, can an object token representing a particular, individual dog be both an instance of the entry-level category [dog], as well as more abstract categories such as [mammal] or [animal], while at the same time being represented as both an individual entity with mereological complexity at multiple scales,

*See also https://chrisfieldsresearch.com/cats.htm for a less formal review of the methods and tools used here.

and as a proper component of even more complex entities? As both abstraction and mereology contribute to the construction of prior probabilities, and to the regulation of precision or attention within Bayesian classifiers (Friston, 2010), the question of how these representations interact – from an neural implementation perspective, how they cross-modulate each other – is crucial to understanding both how mereologically complex objects are identified as individual entities, and how identifiable individual entities are recognized as being mereologically complex.

Our interests here are neither metaphysical nor normative, but rather concern the emprically-characterized functional architecture of categorization and object tracking over time implemented by humans and the potential utility of Chu space and Channel-theoretic methods to describe this functional architecture formally. We make no particular assumptions about the ontological structure of the world or the objects being perceived or about the reality or lack thereof of "universals" or "particulars" as ontological entities, and make no claims about the veridicality or otherwise of human perception. See Churchland (2012) for a treatment of these ontological questions along realist lines that assumes object persistence as an ontological fact, and Hoffman, Singh and Prakash (2015) for arguments that such realist assumptions are inconsistent with an evolutionary account of the architecture of perceptual systems.

We begin this second part of our two-part work in §2 with brief reviews of human perception, categorization and attention as neurocognitive processes and of multi-layer recurrent networks as models of these processes (e.g. Friston, 2010; Grossberg, 2013; Spratling, 2016). In §3, we redescribe perception and categorization, using the Chu space and Channel Theory tools assembled in Fields and Glazebrook (2018), in a way that makes explicit the dualities between dynamic and static properties, individuals and categories, and states and events. We capture these dualities in a "Cone-Cocone Diagram" that formalizes the inferential steps required to link object tokens together to produce a "history" of a persistent object. Such a diagram is seen to be structured by networks of logic infomorphisms and ontologies, as developed in Fields and Glazebrook (2018, §7-§8), and is naturally a simplicial complex in its own right. We then turn our attention, in §4, to mereological categorization and to the key mereotopological question of how the *boundaries* between the components of a complex object are defined. As scenes are themselves mereological complexes, the boundary construction required for object segmentation emerges as the simplest case of inter-object boundary definition. We conclude by listing some open problems for which the current work provides additional motivation.

2 Perception, categorization and attention as neurocognitive processes

The segmentation of visual scenes into "objects" that are represented as located in space, persistent in time, and having some collection of invariant, identifying properties is fundamental to human cognition. The implementation of this process has been extensively studied in both humans and other mammals for several decades and is known to involve both bottom-up and top-down pathways, although the extent to which object- or higher-level expectations modulate low-level image processing remains controversial (a current discussion appears in Firestone and Scholl, 2016). Object identification and categorization depends strongly on context (see Wagemans et al., 2012a,b, for recent reviews of low- and higher-level effects, respectively) and laboratory experiments are typically conducted so as to minimize such context effects (e.g. Brady et al., 2008). Visual object identification is generally conceptualized within a realist ontology in which the perceived objects have observer-independent locations and features, but this is not necessary (Fields et al., 2018). Here we focus on the data structures and processing employed in object identification and categorization – what Marr (1982) called the algorithmic/representational level of analysis and Pylyshyn (1984) termed the functional architecture – with some pointers to the relevant implementation-level neuroscience (see Fields, 2013, for a review of implementation details).

2.1 Dual-pathway vision and object files

The primate, and in particular human, visual system comprises two early processing pathways, a dorsal pathway specialized for the rapid processing of location and motion information, and a ventral pathway specialized for static (e.g. shape, size, texture and colour) feature information (reviewed by Goodale and Milner, 1992; Flombaum, Scholl and Santos, 2008; Fields, 2011); see Cloutman (2013) for a discussion of interactions between these pathways, and Alain et al. (2001); Sathian et al. (2011) for evidence that auditory and haptic perception, respectively, have a similar dual-stream organization. Perception of a located, featured object requires processing by both pathways followed by fusion of the intermediate partial representations they produce.

Studies of visual object tracking over short (500 ms to a few seconds) time periods consistently show that trajectory information dominates static feature information in determining object identity (Flombaum, Scholl and Santos, 2008; Fields, 2011). Kahneman, Triesman and Gibbs (1992) termed the initial, transient representation of a moving object in visual short-term memory the "object file" (see also Treisman, 2006). As under ordinary circumstances all objects are effectively moving due to visual saccades, object files are at least typically initiated by dorsal-stream processing. Static feature information extracted from the relevant part of the visual field by ventral-stream processing is then bound to this initially motion-based representation. These processing steps require 50 – 100 ms in humans, much shorter than the time required for reportable visual awareness of the object.

The object file is the fundamental "token" representation of a located, bounded, featured entity that is distinguished from the "background" of a visual scene. It represents where the object is, its visually-identifiable features, and its instantaneous trajectory during the time window Δt from object-file initiation to feature binding. All further information about the object is added by downstream processing; in particular, whether the object is novel or something previously encountered, either as a type or as a specific individual, must be computed from information available in memory.

With a few prominent exceptions, e.g. LIDA (Franklin and Patterson, 2006; Franklin et al., 2012) and MACSi (Nguyen et al., 2013; Lyubova, Ivaldi and Filliat, 2016), standard computer vision architectures focus on feature extraction and feature-based as opposed to motion-based object identification (see Kotseruba and Tsotsos, 2018, for review); hence they do not directly replicate the human dual-stream process or object-file data structure. Whether doing so provides significant advantages in AI systems remains to be determined.

2.2 Feature-category binding and object tokens

Object files are implemented in a content-dependent way across the posterior temporal cortex (Martin, 2007; Mahon and Caramazza, 2009; Fields, 2013); features of entities perceived as agents or non-agents, for example, are encoded in the lateral or medial, respectively, fusiform gyrus (Fig.

1). This distributed, content-dependent encoding indicates that top-down category information, e.g. agent versus non-agent, is already active in the binding of location and motion information to feature information at the level of the object file.

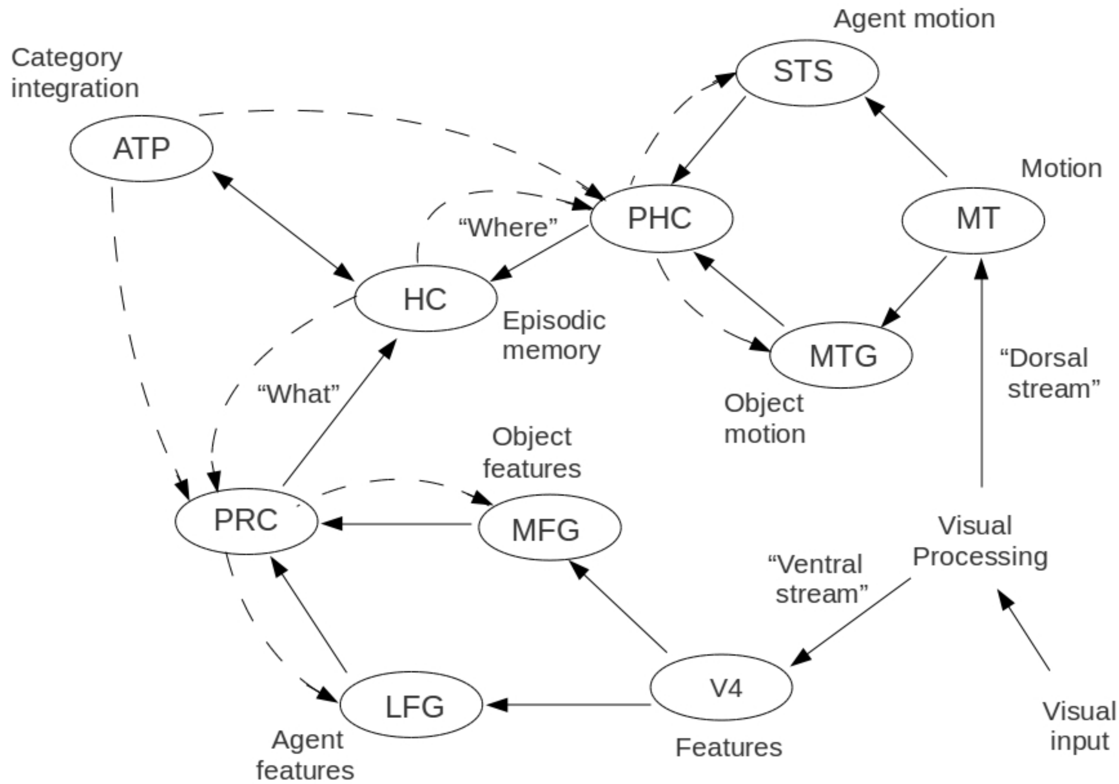

Fig. 1: Simplified functional architecture of human visual object perception within the temporal lobe. Abbreviations are: MT, medial temporal area; STS, superior temporal sulcus; MTG, medial temporal gyrus; PHC, parahippocampal cortex; V4, visual area 4 (occipital cortex); MFG, medial fusiform gyrus; LFG, lateral fusiform gyrus; PRC, perirhinal cortex; HC, hippocampus; ATP, anterior temporal pole. Solid lines are feedforward; dashed lines feedback. Adapted from Fields (2013).

The functional architecture supporting object representation and object-directed attention is already present at birth and its functionality rapidly matures toward adult levels during the first two years; see e.g. Gao et al. (2015); Huang et al (2015) for neuroarchitectural and Johnson and Hannon (2015) for behavioural evidence. Visual feature identification, segregation of co-moving, conjoined objects from other objects and the background, and the complementary process of grouping co-moving, non-conjoined objects, e.g. point-light walkers (Blake and Shiffrar, 2007), are highly dependent on top-down, memory-driven categorization or, in Bayesian terms, expectation confirmation or disconfirmation. Four- to six-month old human infants, for example, typically do not segregate static or co-moving conjoined objects that older infants, children or adults do segregate, but quickly learn to do so when the objects are separately manipulated (Johnson and Hannon, 2015). Young infants similarly fail to group co-moving, non-conjoined objects (i.e. fail to perform "object completion") that older infants, children or adults do group, with the exception of point-light walkers exhibiting biological motion, which infants perceive as single entities from the earliest ages tested (Johnson and Hannon, 2015); see Schlesinger et al. (2012) for a replication of a canonical object completion experiment in the iCub robot. Object categorization is already robust in late infancy (Rakison and Yermoleva, 2010) and develops rapidly with the onset of spoken language; however, language acquisition is not a requirement for functional categorization ability in otherwise cognitively-typical adults (e.g. Schaller, 2012).

Young infants exhibit robust object memory, particularly for familiar faces, and emotional responses to objects, again from the earliest ages tested. Feelings of familiarity and their attendant emotions correlate with feature-based object recognition at the level of perirhinal cortex (Eichenbaum, Yonelinas and Ranganath, 2007). Memory for a particular, individual, re-identifiable object requires a memory-resident representation of that individual object, what Zimmer and Ecker (2010) have termed an "object token." Recognizing a novel object as a distinct, individual thing not encountered before involves encoding a new object token specifically for it. Recognition or re-identification of the same individual object on a later occassion is then a process of matching the current object file to this previously-encoded object token (Fig. 2). This process is, in general, not straightforward, as object features, behaviours, and locations may change between encounters. Even very young infants expect identified objects to maintain constant features, behaviours, and locations over periods of non-observation of seconds to a few minutes (Baillargeon, 2008; Baillargeon et al., 2012). Both feature matching and, after about four years of age, the construction of unobserved and hence confabulated, i.e. fictive causal histories (FCHs) of objects are employed to establish individual object identity across observations separated by more than a few minutes (reviewed in Fields, 2012). Enabling object recognition across feature, behaviour, and context changes requires object tokens to have a "core" of essential properties that change only slowly through time. The distinction between core and variable properties in object tokens is category-dependent and not well understood (see Scholl, 2007, for review).

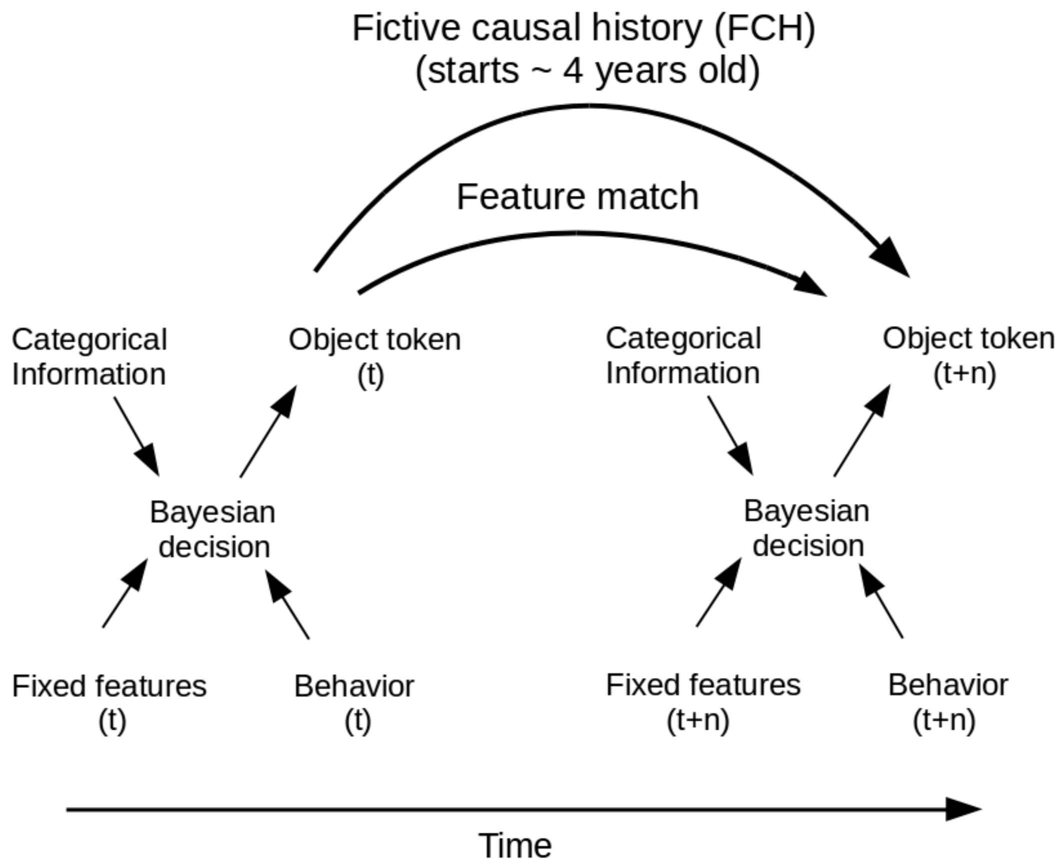

Fig. 2: Identifying an object as the same individual across time requires matching a current object file to a memory-resident object token. Both feature matching and the construction of unobserved (fictive) causal histories (FCHs) are employed to link object tokens across observations (Fields, 2012).

Eichenbaum, Yonelinas and Ranganath (2007) emphasize that categorization by type precedes object-token matching; e.g. an individual person is typically recognized as a person before they are identified as a particular individual person. The feeling of familiarity as well as other emotional responses are generated already at the level of type recognition; Eichenbaum, Yonelinas and Ranganath (2007) consider the example of recognizing a particular person both as a person and as a specific individual person encountered before, but only after some deliberation as a known, named, individual person represented by a memory-resident object token. How categories are related to object tokens in semantic memory, e.g. how the category [person] is related to object tokens representing particular, known persons is not well understood; in particular, whether categories are encoded as prototypes with links to exemplars or as sets of exemplars from which prototypes can be generated remains controversial (see Hampton, 2016, for review). How the object token - to - category relationship is implemented at the neural circuit level, i.e. the details of the circuitry connecting ATP and PRC in Fig. 1, is also not well understood (see Martin, 2007; Keifer and Pulvermüller, 2012, for reviews). While both categorization and individual recognition are generally assumed to be implemented by some form of Bayesian predictive coding (Friston, 2010; Maloney and Zhang, 2010), entry-level categorization is known to occur first (Clarke and Tyler, 2015), and exemplars similar to category prototypes are recognized more easily than less-similar exemplars (Winkielman et al., 2006), how categorization constrains or guides object token matching in particular cases is also not well characterized. The category - to - object token satisfaction relation $\Vdash$ remains, in other words, an empirical question for each particular type-token pair.

2.3 Context perception, event files and episodic memories

Objects are invariably recognized in some context, typically one involving other recognized or at least categorized objects. The spatial "where" information processed by the dorsal stream (cf. Fig. 1) provides the "container" for this context as well as the relative locations and motions of objects within it. Contexts typically, however, also include "how" and "why" information, largely derived by processing pathways in parietal cortex (Fields, 2013), that represent inferences about mechanical and intentional causation, respectively. As in the case of object categorization and individual object-token encoding, these causal inference capabilities are present in rudimentary form in early infancy, and develop rapidly over the first two years (Baillargeon et al., 2012; Johnson and Hannon, 2015). Context assembly has been mapped to parahippocampal cortex, with object token to context binding implemented by the hippocampus (Eichenbaum, Yonelinas and Ranganath, 2007; Ranganath, 2010; Fields, 2013; Ritchey, Libby and Ranganath, 2015). Hommel (2004) has termed the fully-bound representation of interacting objects in context an "event file"; these representations mediate event understanding and context-dependent action planning. Event files are the least complex visual representations that typically enter human awareness; hence they can be considered to be implemented by coherent activity at the level of the global neuronal workspace (GNW, cf. Fields and Glazebrook, 2018, Remark 7.1). For specific GNW-based models of event awareness, see Dehaene and Naccache (2001); Baars and Franklin (2003); Dehaene and Changeux (2004); Baars, Franklin and Ramsoy (2013); Dehaene, Charles and King (2014); see also Franklin and Patterson (2006); Franklin et al. (2012) for discussions GNW design principles as providing the motivation for LIDA.

Event files correspond to "episodes" in episodic memory, again a hippocampus-centered function (Eichenbaum, Yonelinas and Ranganath, 2007; Rugg and Vilberg, 2013; Ritchey, Libby and Ranganath, 2015). As sequences of episodic memories typically contain many of the same "players" – including in particular the remembering subject represented as a "self" (Renoult et al., 2012) – they pose a particular problem for object-token updating. Each episodic memory must contain at least some episode-specific details, e.g. what a particular person was wearing, in addition to the "core" identifying information for each included object token. Linking episodic memories into a temporal sequence requires maintaining this "core" of each object token – which it is useful to consider as a "singular category" with just one member – while modifying the "essential" identifying information for the represented object as needed, e.g. updating a person's age or personality characteristics (Fig. 3). This maintenance process is effectively the construction of a history or causal model of the individual represented by the object token. As such histories occur between episodic memories and are by definition unobserved, they must be confabulated; they are FCHs as discussed above (Fields, 2012). Episodic memory recall and reconsolidation can modify the properties associated with or even the presence of the object tokens referenced by the memory, demonstrating the fragility of such FCHs (Schwabe, Nader and Pruessner, 2014). Infants are capable of episodic recall over short periods – e.g. the time periods required for experiments assaying causal inference – but have limited recall and reconsolidation ability over longer periods (Hayne, 2004; Bauer, 2006). Hence infants can be expected not to maintain robust object histories until about age four, the age when FCH construction typically starts and "childhood amnesia" typically ends.

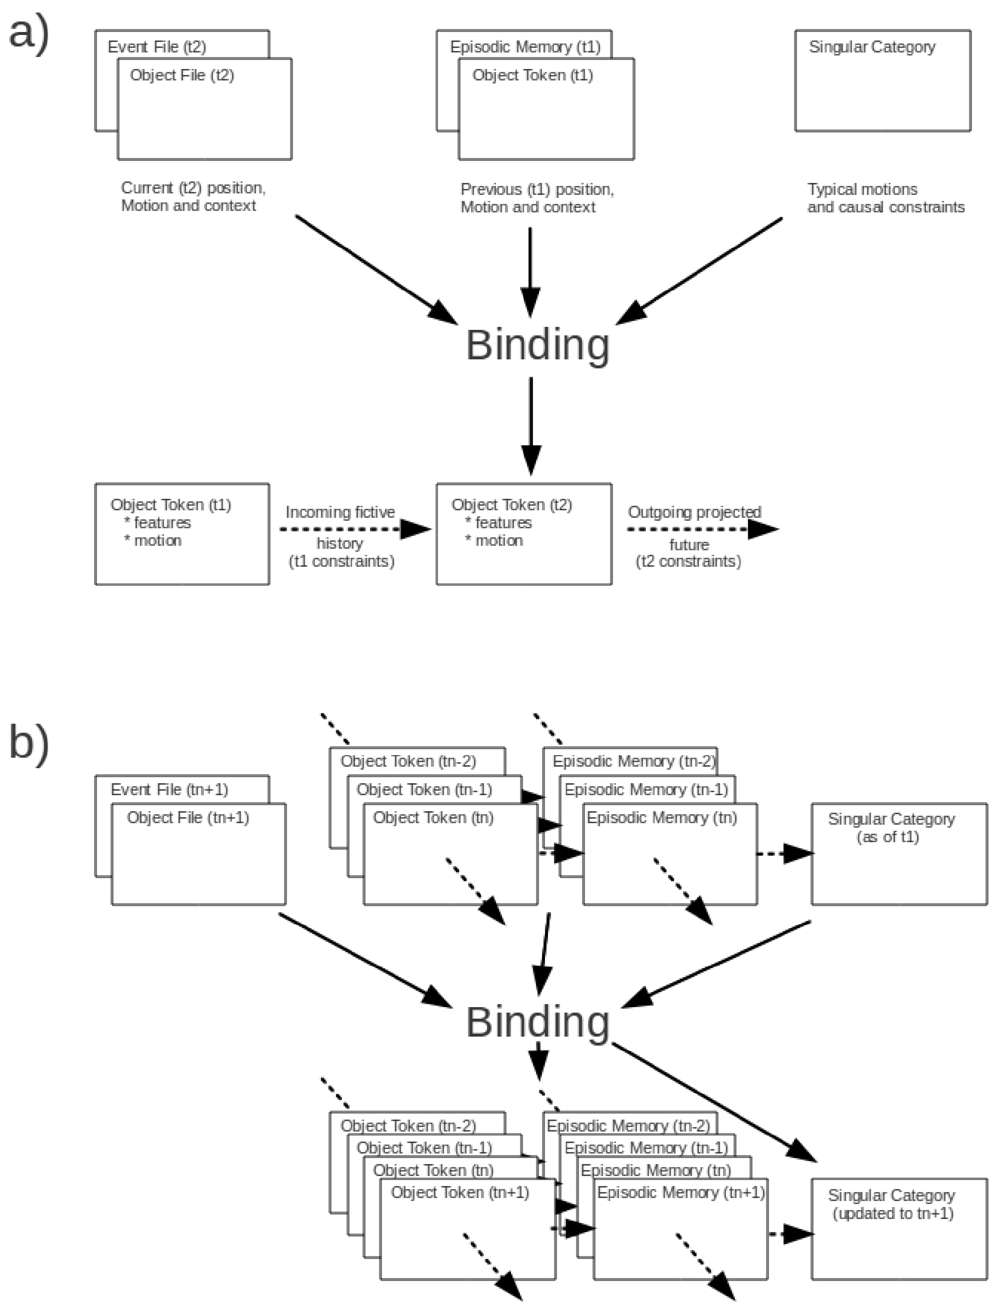

Fig. 3: a) Object token updating following a new event involving a recognized ob-

ject. Successful binding of a current object file to a memory-resident object token that is linked to an episodic memory requires the construction of an FCH that "explains" how the current object could have gotten to the current context from the remembered one. The object token constructed to represent the current object in the current context incorporates updated "non-essential" features as well as the current location and motion information. The singular category specifying the object's "core" identifying criteria constrains the recognition and updating processes. b) Updating the constraint information in a object-specifying singular category defining the "essential" or "core" characteristics of an individual object given a sequence of episodic memories and a new event. Such "core" updates, e.g. of a person's apparent age or personality characteristics, must be infrequent compared to object token updates to maintain a coherent individual object history and hence coherent identification criteria. Adapted from Fields (2012).

Episodic memories can also reference object tokens for individuated and categorized but otherwise unidentified objects, e.g. "some other people" present at a meeting or "other cars" involved in an accident. These "other" objects may appear in no other episodic memories and have no associated histories; they are represented by effectively one-off object tokens that are required by the data structure but play no other role in the system. The human ability to learn to recognize new individuals indicates that such minimally filled-out object tokens are available for matching new incoming object files; however, their lifetime and the extent and context-dependence of their availability remain poorly understood.

2.4 Attention, salience and Bayesian precision

Systems with limited cognitive resources must allocate processing to the inputs most likely to be important. In the current setting, this corresponds to paying attention to some objects and not others. Attentional control in primates is implemented by competing, cross-modulating dorsal (top-down, goal-driven, proactive) and ventral (bottom-up, percept-driven, reactive) attention systems (Goodale, 2014; Vossel, Geng and Fink, 2014). The "salience network" that controls these attention systems develops in concert with the medial-temporal object recognition network, starting from earliest infancy (Gao et al., 2015; Uddin, 2015).

Baysian predictive coding has long been employed as a model of perceptual processing from early vision through categorization and individual object identification (Friston, 2010; Maloney and Zhang, 2010; Spratling, 2016). Bastos et al. (2012) review structural and functional evidence that predictive coding is implemented at the level of local microcircuits comprising cortical mini-columns, the dominant architectural units in mammalian cortex, as well as at the larger scales of functional networks responsible for trajectory recognition, categorization or object-token matching. More recently, Spratling (2017) has implemented a multi-scale predictive coding model of object recognition and demonstrated its efficacy for objects in natural scenes. This use of the same or similar processing methods at different scales, the overall quasi-hierarchical organization of perceptual processing (Van Essen, Anderson and Felleman, 1992), and the small-world, i.e. scale-free structure observed in functional-imaging studies (Rubinov and Sporns, 2010; Sporns, 2013) all suggest that visual object identification exhibits the kind of association between scale and coarse-graining introduced in Fields and Glazebrook (2018, §4-§5). The sections that follow employ the category-theoretic concepts and tools developed in Fields and Glazebrook (2018) to examine and refine this idea that the neurocognitive architecture of human visual object identification may be both structurally and functionally scale-free.

In a Bayesian predictive coding system, attentional control can be modelled by varying the precision assigned to inputs and expectations. In the Bayesian "active inference" framework of Friston (2010), relatively high-precision inputs drive the revision of expectations and model reactive, ventral attention, while relatively high-precision expectations drive input-changing behaviour and model proactive dorsal attention. Attention-switching in this framework is dependent on overall cognitive-affective state, with affective inputs regulating approach and avoidance particularly critical for inducing reactive attention. The framework also allows direct alterations of precision assignments as inferential outcomes (Friston et al., 2015). The adaptive resonance (ART) framework of Grossberg (2013) provides a functionally similar model of attentional control, although its motivation and underlying ideas are distinct from those of Friston (2010) and its state-updating rules are not Bayesian. When viewed as implementations of constraint hierarchies, however, both Bayesian active inference and ART exhibit the deep duality between bottom-up and top-down information flows characterized formally in §3.4 below. It is this duality, we will argue, that makes them useful models of attentional control.

3 Tokens, types and information flow in perception and categorization

While the human classification of perceived objects into cognitive categories corresponding to verbally-expressible concepts like "person" or "house" forms part of the motivation for the work of Dretske (1981); Barwise and Perry (1983); Barwise and Seligman (1997) and others reviewed in Fields and Glazebrook (2018), category-theoretic methods have yet to be applied to the analysis of these processes at the level of detail reviewed in §2 above. The formal treatment of ontologies discussed in Fields and Glazebrook (2018, §7.3), for example, does not explicitly address the question of how ontologies are constructed or maintained through time in the face of new observations. While models at this level of abstraction are of interest conceptually, they are not sufficiently detailed to generate testable predictions. We begin in this section to develop the constructs needed to build models at the level of detail needed to make contact with the empirical results reviewed in §2. We introduce, in particular, the idea of a *Cone-Cocone Diagram* (CCCD) to capture the deep duality evident in the bidirectional flow of constraints between perception and categorization. We point out ways in which "natural" category-theoretic concepts, particularly duality, help to make sense of empirical findings, and identify open issues for either formal or empirical investigation.

3.1 Representing object files in a Chu space

The fundamental perceptual token, the initial representation of a discrete perceptual entity, is the object file. As outlined in §2.1 above, an object file binds a collection of static features such as size, shape, texture and colour extracted by ventral-stream processing to "instantaneous" (i.e. within the Δt of visual short-term memory) location and trajectory information extracted by dorsal stream processing. Recall from Fields and Glazebrook (2018, §2.1) that a Chu space is just a set of objects and attributes organized by a satisfaction relation $\Vdash$. Let $F_1 \dots F_n$ be a finite tuple of static features, each of which can have any one of m distinct values; e.g. if F_i is 'colour' its distinct values are the colours distinguishable, at its finite resolution, by the visual system of interest.

Assume for simplicity that each of these values characterizes the object uniformly. We can then consider a finite binary array $F = [f_{ij}]$, where $f_{ij} = 1$ for some object if and only if feature F_i of that object has its j^{th} possible value. We can similarly consider finite binary arrays $X = [x_{ijk}]$ of discrete instantaneous three-dimensional locations and $V = [v_{ijk}]$ of discrete instantaneous three-dimensional velocities. We will restrict attention to the case in which every object has some value for every perceptible feature, a single instantaneous location, and a single instantaneous velocity; feature changes and extended curvilinear trajectories are treated with sequences of such simplified files. In this case, we can characterize an *object file* as an instance of the finite array $[F, X, V]$. The set of all possible instances of $[F, X, V]$, and hence the set of all possible object files that can be generated by the visual system of interest, is a finite set $\{O_i\}$. This set becomes larger, but still remains finite, when complications such as non-uniform features or curvilinear trajectories are included in the description, provided that the resolution of the system remains finite on all dimensions.

The most fundamental abstraction implemented by the visual system is *object permanence*, i.e. the maintenance of object identity over time (see Flombaum, Scholl and Santos, 2008; Fields, 2011, 2017, for reviews). At the level of the object file, the relevant timeframe for object permanence is a "view" lasting between half a second and a few seconds. Objects that remain fully or even partially visible during such a view are considered to remain "the same thing" while seen. Whether an object that does not remain visible is perceived as remaining "the same thing" during a view depends on the age of the perceiver (less or more than 1 year) and the details of its occluded trajectory (Fields, 2011). Objects moving sufficiently fast are "seen" as persistent even if their static features, e.g. size, shape or colour, vary over considerable ranges (Flombaum, Scholl and Santos, 2008). The relative insensitivity to static features of short-term object permanence indicates that it is a primarily dorsal-stream phenomenon.

Let $\{C_i\}$ be the finite set of finite (indeed short) sequences of object files that are treated by the cognitive system of interest as indicating object permanence during the course of a single view. The elements of $\{C_i\}$ are then natural "types" relative to the "tokens" in the set $\{O_i\}$ of possible object files; an element $C_i \in \{C_i\}$ can be though of as "associating" a sequence of object files into a single abstracted representation. Hence we can consider

$$\mathcal{C}_i = (\{O_i\}, \Vdash_P, \{C_i\}) \tag{3.1}$$

to be a Chu space, where here $\Vdash_P$ is the empirically-determined relation "consistent with object permanence" defined on sequences of object files. As noted above, this $\Vdash_P$ is strongly dependent on trajectory and occlusion but relatively independent of static feature constancy. The Chu space $\mathcal{C}_i$ clearly describes a Classification in the sense of Barwise and Seligman (1997) as discussed in Fields and Glazebrook (2018, §6.1). Note that an element of $\{C_i\}$ may not be a concept in the sense of Fields and Glazebrook (2018, §3), as the value of every feature as well as the position and velocity can, at least in principle, change between every object file contributing to the perception of a persistent object. The development of object trackers implementing a relation $\Vdash_P$ approximating that of human vision remains a substantial AI challenge (Kristan et al., 2015).

The association of object files into an element of $\{C_i\}$ adds top-down, expectation-based information about *identity over time* to the "raw" information of perception. This added information may, in fact, be incorrect; trajectories that appear to preserve object identity may involve distinct objects, while those that appear not to preserve object identity may involve a single object (Fields, 2011). More subtly, association of object files into an element of $\{C_i\}$ also subtracts information by suppressing motion information available at the individual object-file level relative to shared or average feature information. While trajectory information dominates feature information in determining which sequences of object files to treat as indicating persistence and hence to associate, persistent objects are required (by humans) to have persistent features, at least during the course of short, single-view interactions (Baillargeon, 2008). Conferring persistence on an object converts its observed motion into an abstracted, categorizable "behaviour" that the object may or may not execute on other occasions. We will see this combination of top-down information addition (abstraction to an invariant, here identity over time) and bottom-up information loss (coarse-graining, here of both motion and feature information) repeated at multiple scales. We develop a formal representation of this process in §3.4 below.

3.2 From object files to object tokens and object histories

Persistent objects are the "entities" in the common-sense ontology humans typically develop in late infancy. These entities participate in episodic memories and are represented by object tokens and, if they recur sufficiently often to be recognized as persistent individuals, by singular categories and (largely fictive) histories. As with the initial abstraction of persistence, these successive levels of abstraction both add and subtract information. Types at one level of abstraction, in particular, become tokens at the next.

As representations of persistent objects, *object tokens* can be identified with elements of the set $\{C_i\}$, with the abstract "type" representations of finite sequences of object tokens. The functional context in which object tokens are active is not only more coarse-grained than that of object files; it implements a distinct set of constraints. While object files represent motion at high resolution, motion information is suppressed at the object token stage to enable re-identification of individual objects regardless of how they are moving Fields (2011). Object tokens are, at the time of their construction, already represented as instances of multiple types (Fig. 4). All persistent objects are instances, first, of the types representing their visually-identified features. They are, second, classified automatically by threat detection, agency detection and animacy detection systems active beginning in early infancy (Fields, 2014); the presence of a face alone indicates agency to human infants. They are also classified, when possible, into entry-level and then more abstract cognitive categories, an ability also developed in infancy (Rakison and Yermoleva, 2010). These token - type relationships can be represented as Classifications, as is standard in the literature (e.g. Barwise and Seligman, 1997), and as surveyed in Fields and Glazebrook (2018). Recognition of an object by type generates a feeling of familiarity with the *type*; e.g. seeing a cat generates a feeling of familiarity with cats (Eichenbaum, Yonelinas and Ranganath, 2007).

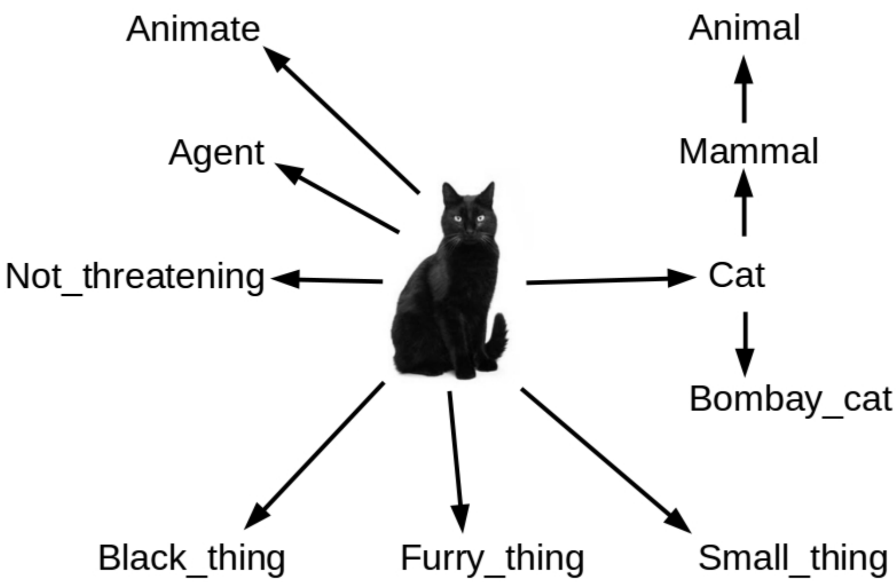

Fig. 4: An object token is classified at construction into multiple types by distinct but cross-modulating processes. These include animacy and agency detection, emotion-mediated threat detection, and entry-level followed by superordinate and subordinate categorization into "types" of object.

Here we are primarily interested in the *re-identification* of individual objects, i.e. the creation of an association indicating identity, and hence persistence over time, between an object token constructed now and one constructed previously. At the object token level, the relevant timeframes for persistence range from the few seconds separating views to the decades separating a high-school graduation from a 50^{th} reunion. How humans are capable accurately re-identifying individual objects across such long gaps in observation, and despite the static feature, behaviour, and context changes such long gaps typically entail, has posed a problem to philosophers for millennia (Scholl, 2007) and remains a critical question for both cognitive neuroscience and AI (Fields, 2016).

Let $C_i(t_1), C_j(t_2), ...C_k(t_n)$ be a sequence of n object tokens encountered at successive times, possibly with long gaps between observations. Recognizing successive object tokens as tokens of the very same individual thing involves at least the two processes discussed in §2.3 above, i.e. matching to a set of core features composing a singular category and linking via FCH construction, with the uncertainty associated with each process increasing with the time between perceptual encounters. Let $D_l[t_1, t_n]$ comprise both the singular category and the FCH that together confer persistence on the object-token sequence $C_i(t_1), C_j(t_2), ...C_k(t_n)$, and let $\{D_l[t_1, t_n]\}$ be the set of all such representations over sequences of elements of $\{C_i\}$ indexed by observation times in the closed interval $[t_1, t_n]$. Each element of $\{D_l[t_1, t_n]\}$ represents the "package" of evidence used by the perceptual agent to identify some persistent object l between times t_1 and t_n; the entire set $\{D_l[t_1, t_n]\}$ specifies the identification criteria for all of the objects the agent is capable of identifying during this time interval. As in the case of sequences of object files discussed in §3.1 above, these elements are natural "types" for the object tokens over which they are defined. Hence we can consider a Chu space or Classification (in the sense of Fields and Glazebrook (2018, §6.1)) $\mathcal{A}_i[t_1, t_n]$ as given by:

$$\mathcal{A}_i[t_1, t_n] := \langle \{C_i[t_1, t_n]\}, \{D_i[t_1, t_n]\}, \Vdash_P [t_1, t_n] \rangle \tag{3.2}$$

where $\Vdash_P [t_1, t_n]$ is the empirically-determined relation "consistent with object permanence" defined on sequences of object tokens between t_1 and t_n. As with the short-term persistence relation $\Vdash_P$ between object files defined in §3.1, determining this longer-term persistence relation $\Vdash_P [t_1, t_n]$ for particular human subjects and particular objects is a difficult empirical question (see e.g. Nichols and Bruno, 2010, for experiments designed to determine the criteria people use to identify other people as persistent over time). Implementing this relation requires solving an instance of the frame problem, i.e. determining the set of inter-context changes that do not alter the identity of an object even though they may radically alter its features or behaviour (Fields, 2013). The continuing efforts to develop facial-recognition systems robust against everyday feature and context changes illustrate the difficulty of this problem (Patel, Kothari and Bhurchandi, 2015).

The set of time-indexed representations $\{D_l[t_1, t_n]\}$ can be conceptualized more abstractly by noting that at each t_j, the set of possible object tokens $\{C_i(t_j)\}$ is also the set of types of a classification. For each single time step $t_j \longrightarrow t_k$, the persistence criterion $\Vdash_P [t_j, t_k]$ induces maps – what we have called FCHs – between pairs of object tokens that can be consistently considered to be tokens of the same individual object (Fig. 5a). These FCHs, together with the maps (here assumed to be identities) linking the singular categories for persistent objects, can be considered infomorphisms between the underlying classifications at t_j and t_k. It is then natural to interpret the set $\{D_l[t_j, t_k]\}$ as a channel between the underlying classifications; this channel comprises, intuitively, the (assumed constant) singular categories and the constructed FCHs (Fig. 5b). Extending the process of linking object tokens by FCHs forward in time results in a hierarchy of channels, with the most temporally-extended channel as the colimit (Fig. 5c). Dropping the explicit $\{\}$ to simplify the notation, the colimit cocone $D_i[t_1, t_n]$ admits a vertex classification, which we denote $\mathbf{C}_i$ (with time interval $[t_1, t_n]$ understood). Recall that this $\mathbf{C}_i$ is induced by a complex of infomorphisms:

$$\cdots \longrightarrow \mathcal{A}_i[t_1, t_n] \longrightarrow \mathcal{A}_{i+1}[t_1, t_n] \longrightarrow \cdots \tag{3.3}$$

as depicted in Fields and Glazebrook (2018, (6.7)). We will refer to such diagrams $D_i[t_1, t_n]$ as "Cocone Diagrams" or CCDs extending for a specified time interval, e.g. $t_1...t_n$ in Fig. 5c.

a)

$$\begin{tikzcd}
\mathrm{SC_i(t_j)} \arrow[r, "\mathrm{Id}"] & \mathrm{SC_i(t_k)} \\
\mathrm{OT_i(t_j)} \arrow[u, "\mathrm{Inst}"] & \mathrm{OT_i(t_k)} \arrow[u, "\mathrm{Inst}"'] \arrow[l, "\mathrm{FCH}"]
\end{tikzcd}$$

b)

$$
\begin{tikzcd}
& D_i[t_1, t_2] & & & D_i[t_3, t_4] & \\
C_i(t_1) \arrow[ur] & & C_i(t_2) \arrow[ul] \arrow[ll, "\text{FCH}"'] & C_i(t_3) \arrow[ur] & & C_i(t_4) \arrow[ul] \arrow[ll, "\text{FCH}"']
\end{tikzcd}
$$

c)

$$
\begin{tikzcd}
& & & D_i[t_1, t_n] & & & \\
& & D_i[t_1, t_4] \arrow[ur, dotted] & \cdots & D_i[t_{n-3}, t_n] \arrow[ul, dotted] & & \\
& D_i[t_1, t_2] \arrow[ur, dotted] \arrow[r, dotted] & & D_i[t_3, t_4] \arrow[ul, dotted] \arrow[r, dotted] & & D_i[t_{n-1}, t_n] \arrow[ul, dotted] & \\
C_i(t_1) \arrow[ur] & C_i(t_2) \arrow[u] \arrow[l, "\text{FCH}"'] & C_i(t_3) \arrow[ur] \arrow[l, dotted] & C_i(t_4) \arrow[u] \arrow[l, "\text{FCH}"'] & & C_i(t_{n-1}) \arrow[u] \arrow[l, dotted] & C_i(t_n) \arrow[ul] \arrow[l, "\text{FCH}"']
\end{tikzcd}
$$

Fig. 5: a) Interpreting sequential object tokens as representing the same persistent individual constructs an FCH to link them. The FCH is depicted as acting backwards in time as it is built from the new observation to the old one. Here $\mathrm{SC_i(t_j)}$ denotes Singular Category i at t_j, $\mathrm{OT_i(t_j)}$ denotes Object Token i at t_j, etc. Id = Identity, and Inst = Instance. b) Families of FCHs link sets $C_i(t_j)$ of object tokens instantiated at different times. The set $D_i[t_j, t_k]$ of abstracted singular category plus FCH pairs representing objects persistent from t_j to t_k can be viewed as a channel between sets of linked object tokens. c) A CCD representing an object history during a time interval $t_1 ... t_n$. The set $D_i[t_1, t_n]$ is a colimit cocone for sequences of object tokens consistent with persistence from t_1 to t_n.

The extension of Fig. 5b to the colimit, Fig. 5c embodies a theoretical prediction, *viz.* that histories representing individual entities are extended forward in time without altering their represented "past" states. This prediction sometimes fails, e.g. when an artwork is discovered to be a forgery or when a trusted civil servant is discovered to be a spy. That such failures are exceptional and tend to provoke cognitive crises indicates that the prediction of incremental, forward history construction is a good approximation. Considering the "essential" identifying properties of objects to remain constant is clearly also an approximation; such properties can change over time, though they cannot all change together without causing identification failure. In the case of human beings, for example, both (approximate) age and core personality characteristics are identifying properties (Nichols and Bruno, 2010); hence a child with an adult friend's personality is not identified as one's adult friend. Slow, asynchronous changes in the composition of singular categories and hence small departures from identity of the linking maps between them do not alter the structures of the above diagrams. Such changes do, however, render FCH construction more difficult.

3.3 Contexts, event files and episodic memories

Objects are never encountered in complete isolation; even the most austere psychophysics experiments have a computer screen and the surrounding laboratory as a context. In real-life settings, objects are typically encountered in interacting groups. Object tokens have, therefore, lateral synchronic associations as well as the diachronic links implemented by FCHs. The event-file construct of Hommel (2004) provides a "snapshot" of such associations over the few-second to few-minute timeframes intuitively regarded as single "events." Event files capture interactions between objects as well as their significance and affective consequences for the observer. These kinds of information provide crucial input into the FCH construction processes that allow the objects participating in the event to be identified (Eichenbaum, Yonelinas and Ranganath, 2007; Zimmer and Ecker, 2010; Fields, 2012).

Event files as defined by Hommel (2004) are effectively tokens; each represents a discrete event that can be encoded and then retrieved as a discrete episodic memory. Single events are by definition localized in time and hence cannot be repeated; recalling an event and hence (partially and perhaps inaccurately) reconstructing an event file, in particular, occurs in a current context and itself constitutes a distinct event. It is now clear that the context in which an episodic memory is recalled influences the recalled content, and the context-dependent process of re-encoding ("reconsolidating") an episodic memory for future access can also alter its content (reviewed by Schwabe, Nader and Pruessner, 2014). Such lability challenges the forward construction of object histories depicted in Fig. 5; however, it is generally considered a deficit of human cognition and hence not a goal for AI systems. The recognition of event *types* as such, independently of the objects involved, is not well characterized experimentally. While recognizing instances of events or event sequences completely defined by known sets of rules (e.g. chess games) may be straightforward given a suitably abstract description, recognition of events not defined by known sets of rules, such as incipient stock market crashes (Fievet and Sornette, 2018), is considerably more challenging. It seems reasonable to expect, however, that event tokens (i.e. event files) and event types can be regarded as forming classifications under the action of a satisfaction relation that maps tokens to types (cf. Fields and Glazebrook, 2018, §7.5-§7.7). We consider this further in §4 below in the broader context of mereological complexity and reasoning.

3.4 Learning new categories and Cone-Cocone Diagrams

With high frequency in infancy and childhood but typically reduced frequency thereafter, humans encounter not just individual objects, but object types that they have never encountered before. Humans often learn to recognize such novelties from just one "training" encounter. Understanding how humans achieve such one-shot learning is a major challenge for cognitive neuroscience, just as replicating this ability is a major challenge for machine learning. As shown by Fei-Fei, Fergus and Perona (2006), one-shot learning of a novel category benefits from the use of all known, categorized object tokens as negative examples; this suggests that negative category links as well as positive ones are needed, and that negative analogs of the construction given by (3.2) must be considered. Besides achieving efficient, preferably one-shot learning from exemplars, this problem has (at least) two additional components: rapidly recognizing novelty (i.e. categorization failure) and switching from classification mode to learning mode. As Oudeyer, Baranes and Kaplan (2013) have emphasized, it is *learnable* novelty that must be recognized; otherwise precious resources are wasted on attempts to learn the unlearnable.

Consider a novel object that is easily classified as an instance of a familiar entry-level category: a novel cat, for example. The object is recognized as novel because its object token does not match any existing singular category, cannot be linked to any existing object token by a plausible FCH, or both. Interacting with the object over an extended period (several views, a few minutes) or encountering it again after a short delay allows certain of its properties to be identified as unchanging; in the case of a cat, these may include size, shape, colour pattern, face and voice but not location or behaviour. The (short) sequence of distinct object tokens recorded during such interactions serves, in other words, to associate some properties of the object into a provisional singular category. The principle of association here is, once again, persistence: each successive object token indicates the object as persistent, and the features encoded by object tokens in the sequence are similar enough to be treated as identical.

We can, in this case, consider the distinct feature instances encoded by the distinct object tokens in the sequence to be feature "tokens" and consider the object tokens themselves, which the criterion of persistence identifies as representing one individual object, as jointly defining a "type" that organizes those tokens. The construction employed in §3.1 above can then be employed to construct a classification of these (feature) tokens into these (object) types. This classification is the Chu-space dual of the classification of object tokens by singular categories shown in Fig. 5a.

The requirement of a familiar entry-level category can now be relaxed: suppose that what is encountered is not a novel cat, but an entirely novel animal, perhaps a pangolin or a platypus. In this case, categorization failure for known animal categories indicates that a new category must be learned, while categorization success (we can assume) under [animal] indicates that it is a new *animal* category that must be learned. Non-matching animal categories then provide negative exemplars. However novel the object encountered is, it must have *some* familiar features; if it is recognized as an animal, these would be shared animal features such as an approximate size and shape, particular shape of the face, and aspects of its behaviour. Even very young infants can use features of these kinds to initiate classification and identify novelty (Rakison and Yermoleva, 2010). Placement in any familiar category allows the construction of a singular category as outlined above. Construction of a non-singular category – e.g. [pangolin] – merely requires abstraction, i.e. allowance of inexact matches.

The problem of maintaining a singular category across changes in essential features introduced in §3.2 above can now be seen as a special case of category learning. A singular category is robust against feature changes if the FCHs linking its instances are strong enough that persistence at the object token level can induce persistence at the singular category level. The "flow of association" in this case is the reverse of that depicted in Fig. 5c; the properties composing the singular category are in this case the "tokens" that are held together by the persistent object history as a "type." This dual view of histories helps explain why "revisionary histories" involving forgery or disloyalty discussed in §3.2 are problematic; they introduce inconsistencies at the singular category level that cause failure of the type (history) – token (singular category instance) classification relation.

Reversing the arrows in a CCD (e.g. Fig. 5c) yields a cone, the dual of a cocone. A system capable of both object history construction and its dual, category learning with singular category maintenance as a special case, is thus characterized by a *Cone-Cocone Diagram* (CCCD); such a diagram can be represented by making all of the arrows in a CCD such as Fig. 5c double-headed. Continuing the notation used in Fig. 5c, we denote the corresponding CCCD by $\mathbf{Dg}_i[t_1, t_n]$.

A CCCD captures the simultaneous upward and downward flow of constraints that characterize human visual object identification and, it is reasonable to suppose, other sensory modalities both functionally and neuro-architecturally (Hochstein and Ahissar, 2002). The duality expressed by a CCCD thus includes the functional duality between high-precision expectations and high-precision inputs in an active inference system Friston (2010), and hence the duality between dorsal (active) and ventral (passive) attention systems. It allows us to understand the central paradox of familiarity, that familiarity can confer either high or low salience in a context-dependent way, from a mechanistic perspective; indeed the "switch" between these dual constraint flows appears to be implemented, in humans, by the amygdala - insula - cingulate axis at the core of the salience network (e.g. Uddin, 2015). A CCCD also, however, captures a more subtle duality between processes: it enables object files, object tokens, and object histories to be viewed not as tokens, but as types that organize, respectively, trajectory components, features, and feature-based singular categories into mutually-consistent collections. Representing visual object identification by a CCCD is, therefore, making a strong empirical prediction: that the satisfaction relations $\Vdash_P$ (3.1), $\Vdash_P$ $[t_1, t_n]$ (3.2) and their extension to the singular category – history relation are not just idiosyncratic results of individual learning histories, but also logical links that construct a "possibility space" for the combined perceptual/conceptual system. Such logical constraints correspond, in a Bayesian picture, to very low and very high prior probabilities being effectively clamped to zero and one respectively; hence violations of these constraints would be expected to generate stronger "conflict" signals than predicted by a continuous probability model. The representation of visual object identification by a CCCD predicts, therefore, systematically poor performance on both low- and high-probability judgments under uncertainty, a phenomenon that is robustly observed (reviewed by Kahneman, 2011). It also predicts a larger affective response to both success and failure of object identification than a continuous model. Whether this is correct remains to be determined, but the high affective response to apparent prediction failures observed in autism (e.g. Lawson, Rees and Friston, 2014; Van de Cruys et a., 2014), for example, renders it plausible.

3.5 Local logics embedded in CCCDs

Recall that any classification generates a natural local logic in accordance with Fields and Glazebrook (2018, Definition 6.6), and that Barwise and Seligman (1997, Prop. 12.7) ensures that any local logic defined on a classification can be identified with the local logic generated by the classification (cf. Fields and Glazebrook, 2018, Example 7.1). These ideas can now be applied to the classifications defined above to characterize the categorization and identity maintenance processes in terms of the actions of local logics (see also Kent, 2016).

To begin, we can immediately apply the principle of Fields and Glazebrook (2018, Definition 6.6) to (3.1) relating "instantaneous" object files (tokens) to short sequences of object files (types) indicating object permanence, to obtain a local logic $\mathsf{Lg}(\mathcal{C}_i)$ with regular theory $\mathsf{Th}(\mathcal{C}_i) = (\{C_i\}, \vdash)$. This $\mathsf{Th}(\mathcal{C}_i) = (\{C_i\}, \vdash)$ expresses the effective criteria for short-term object permanence and hence captures, albeit implicitly, an important part of the semantics of "persistent object" for the system it describes. Likewise, in (3.2), we have a local logic $\mathsf{Lg}(\mathcal{A}_i[t_1, t_n])$ with regular theory $\mathsf{Th}(\mathcal{A}_i[t_1, t_n]) = (\{D_i\}[t_1, t_n], \vdash)$ (for each i) that captures the effective criteria for longer-term object permanence and hence additional components of the semantics of "object." In both cases all tokens are normal as defined in Fields and Glazebrook (2018, Definition 6.5). On recalling Fields and Glazebrook (2018, Definition 6.7), we can take as a working principle that temporal sequences of infomorphisms such as (3.3) are also logic infomorphisms satisfying the properties of Barwise and Seligman (1997, 12.3). Accordingly, an underlying semantic structure is built into Fig. 5c, and hence to the ensuing CCCD diagram $\mathbf{Dg}_i[t_1, t_n]$.

Sequents as introduced in Fields and Glazebrook (2018, §6.7) play an interesting role in this development. On the one hand, they are implicitly assumed in the preceding discussion (see also Fields and Glazebrook (2018, §6.8,§7.3). On the other hand, we recall from Fields and Glazebrook (2018, §7.8) that on relaxing the sequent relation $\vdash$ to a conditional probability (see Fields and Glazebrook, 2018, (7.27)), a sequence of logic infomorphisms may function as a chain of Bayesian inferences. This is consistent with the use of Bayesian methods, particularly Bayesian predictive coding, reviewed in §2.4 above, and suggests that the diagrams $\mathbf{Dg}_i[t_1, t_n]$ may be considered as effective carriers of Bayesian inference through sequences of episodic memories (*pace* issues of memory change under reconsolidation and of low- and high-probability clamping), an idea that is extended in §4 below.

As regards ontologies, we recall that both types C_i and their identification criteria $D_i[t_1, t_n]$ are not strictly speaking sets of (formal) concept symbols. We suggest that a weaker sense of ontology is obtainable using the generic relations $\leq, \perp, |$ in Fields and Glazebrook (2018, Definition 7.3). Generally, however, we can acknowledge the viewpoint of Kalfoglou and Schorlemmer (2003), which sees the local logics themselves as characterizing effective ontologies. This derivative sense of ontology is useful for our purposes since ontological partitions into "entities" tend to induce "spatial" boundaries around conceptual and/or perceptual partitions (Smith, 1996). Such induced boundaries can be identified with coarse-grainings and hence with induced coarse-grained geometries, as will be discussed in §4.3 below. An alternative, more abstract approach to ontologies and the associated knowledge representation is described in Spivak and Kent (2012).

4 Parts and wholes: Using Chu spaces and information channels to represent mereological complexity

The Formal Ontology introduced by Husserl (1970) developed a theory of "parts" and "wholes" towards a foundation for mereological (i.e. part – whole) reasoning, a methodology that also has roots in the works of Aristotle, Brentano, Whitehead, and others (as reviewed and developed in Casati and Varzi, 1999; Lando, 2017; Leśniewski, 1929; Simons, 1987; Smith, 1996). Formalizations of mereological reasoning (e.g. Casati and Varzi, 1999; Smith, 1996) have found wide application in geographic information systems (GIS) and formal ontologies for scientific domains. The implementation of mereological reasoning in humans is not, however, well understood; indeed we have been able to find only a single neuroimaging study explicitly comparing mereological and functional classifications (Muehlhaus et al., 2014). As mereological reasoning appears specifically to fail in the "weak central coherence" (WCC) phenotype of autism spectrum conditions (Happé and Frith, 2006), understanding its implementation and its relation to abstraction, which does not consistently fail in WCC, is potentially of clinical relevance.

4.1 Perceptual identification of mereologically-complex objects

The macroscopic objects perceptible by humans are by definition mereologically complex: they have multiple perceptible parts, each of which has further parts, etc. Such objects can, moreover, be assembled into larger complexes, with perceptual scenes being ubiquitous, transient examples. Many such larger complexes are, however, not transient but rather meaningful, persistent objects in their own right. A fundamental challenge posed by human object perception is to understand what mereological complexes are perceived as "whole" objects (and likewise, which ones are not), how object tokens representing such complexes are constructed, and how such object tokens are linked into persistent histories despite changes in the properties and even identities of the "parts" making up the complex.

A specific example of a mereological hierarchy is shown in Fig. 6. Individual human beings, such as author CF, are entry-level (EL) objects and hence are represented by EL object tokens. Human beings are inevitably members of larger complexes, including families, extended families, tribes, ethnic groups, nations, etc. The smaller instances of such complexes (e.g. human nuclear families) can be directly perceived; larger instances may not be perceptible but can be referred to using language, images, and abstract graphics. Hence object tokens can be constructed for such complexes. Object tokens representing "parts" such as CF are naturally linked to object tokens representing complexes, such as CF's family, by "part_of" relations. Such relations similarly link parts of CF to CF. Entry-level objects appear to play a special role in such hierarchies; "part_of" links are transitive both above and below EL objects, but not across EL objects. A part of CF is not a part of CF's family, just as a part of a car is not a part of a fleet of cars.

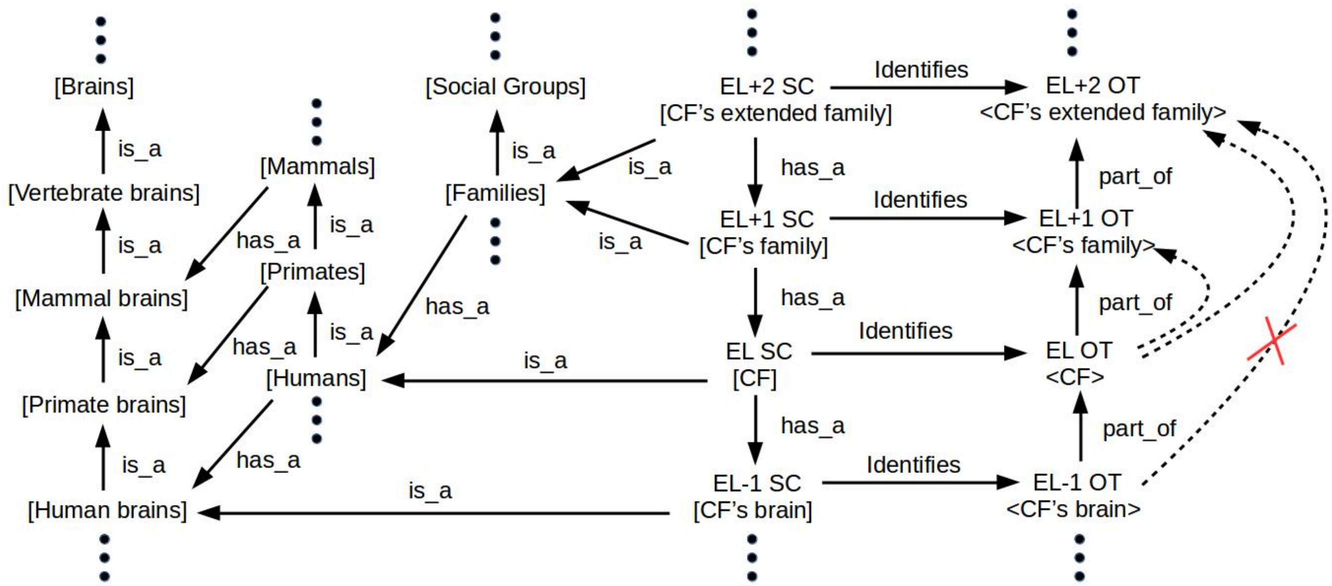

Fig. 6: Example of an object token (OT) hierarchy extending both above and below a mereologically-complex entry-level (EL) object, one of the present authors (CF). Each OT has an associated singular category (SC) specifying identifying static and behavioural features. These SCs are in turn associated with general categories, some of which are shown here. Solid arrows show typical "part_of", "has_a" and "is_a" links. Dashed arrows show induced part_of links; red "X" indicates the failure of "part_of" transitivity across the EL OT.

Hierarchies of tokens linked by "part_of" relations are commonplace in AI systems. Such "part_of" hierarchies raise the questions of what the "part_of" relation is, how it is established, and how it is maintained over time. The correspondence between object tokens and singular categories discussed in §3.2 above provides a partial answer: "part_of" relations between object tokens correspond to "has_a" relations between singular categories, which in turn correspond to "has_a" relations between general categories (Fig. 6). While the object token is the locus of learning for the first exemplars of EL objects encountered in infancy and childhood, once a general category has been learned the "part_of" links between new object tokens can be induced by inter-category "has_a" links. The mechanisms by which such link induction is implemented in humans remain to be elucidated experimentally; we consider formal structures supporting this process below.

4.2 Mereological hierarchies as hierarchies of CCCDs

Let us first consider the Chu space $\mathsf{C} = (C_{\mathsf{o}}, \Vdash_{\mathsf{C}}, C_{\mathsf{a}})$, where C_{o} and C_{a} are sets of object tokens and their corresponding singular categories and $\Vdash_{\mathsf{C}}$ is the "Identifies" relation in Fig. 6. Recall from Fields and Glazebrook (2018, §3.1) the pair of maps (α, ω) (there considered as a Galois connection) given by:

$$\begin{aligned} \alpha &: \mathcal{P}(C_{\mathsf{o}}) \longrightarrow \mathcal{P}(C_{\mathsf{a}}) \text{ with } \alpha(X) = \{\mathsf{a} : \forall x \in X,\ x \Vdash_{\mathsf{C}} \mathsf{a}\} \\ \omega &: \mathcal{P}(C_{\mathsf{a}}) \longrightarrow \mathcal{P}(C_{\mathsf{o}}) \text{ with } \omega(A) = \{x : \forall \mathsf{a} \in A,\ x \Vdash_{\mathsf{C}} \mathsf{a}\}. \end{aligned} \tag{4.1}$$

Here α clearly maps an object token to its (unique) singular category and ω maps a singular category to its (unique) object token. To generalize to the case of objects viewed multiple times, and hence to singular categories linked to sequences of object tokens, we abuse the notation slightly to allow C_{o} to be a set of sequences of object tokens.

Definition 4.1.

(1) Suppose X is a set of subsets consisting of *parts of objects*. Then we define $\omega \circ \alpha(X)$ to be the set of subsets of *whole parts of objects as obtained from* X.

(2) Suppose Y is a set of subsets consisting of *parts of attributes*. Then we define $\alpha \circ \omega(Y)$ to be the set of subsets of *whole parts of attributes as obtained from* Y.

Note that these definitions require the parts of an object (as represented by a sequence of object tokens) to be both persistent and persistently parts of the object; hence they only approximate situations in which an object can lose a part, without replacement, but still maintain its identity. The usage "whole parts" is employed here to emphasize that "wholes" on one level may be "parts" at the level(s) above.

Likewise, for a given classification we have $\mathcal{A} = \langle \mathrm{Tok}(\mathcal{A}), \mathrm{Typ}(\mathcal{A}), \Vdash_{\mathcal{A}} \rangle$, and for $a \in X \subseteq \mathrm{Typ}(\mathcal{A}), \mathrm{b} \in \mathrm{A} \subseteq \mathrm{Tok}(\mathcal{A})$, we have:

$$\begin{aligned} \alpha^* &: \mathcal{P}(\mathrm{Typ}(\mathcal{A})) \longrightarrow \mathcal{P}(\mathrm{Tok}(\mathcal{A})) \text{ with } \alpha^*(\mathrm{X}) = \{\mathrm{b} : \forall \mathrm{a} \in \mathrm{X},\ \mathrm{x} \Vdash_{\mathcal{A}} \mathrm{b}\} \\ \omega^* &: \mathcal{P}(\mathrm{Tok}(\mathcal{A})) \longrightarrow \mathcal{P}(\mathrm{Typ}(\mathcal{A})) \text{ with } \omega^*(\mathrm{A}) = \{\mathrm{a} : \forall \mathrm{b} \in \mathrm{A},\ \mathrm{a} \Vdash_{\mathcal{A}} \mathrm{b}\}. \end{aligned} \tag{4.2}$$

Iterating these conditions allows us to move one rung at a time through the mereological hierarchy when incorporating information channels.

To see how this mereological hierarchy can be constructed, we first of all construct a (quasi-hierarchial) complex of CCCDs following §3.2, §3.4 and §3.5. Recalling (3.2) and drawing on the techniques of Distributed Systems of Barwise and Seligman (1997, §6) (cf. Fields and Glazebrook, 2018, §7.2), we may view the index i as indicating a 'level' in a complex of CCCDs constructed from connected sequences. In such sequences, the time intervals will generally be distinct. Thus we commence with families of time dependent logic infomorphisms arising from such morphisms between different classifications as specified in (3.2):

$$\mathcal{A}_i[t_{i_1}, t_{i_n}] \longrightarrow \mathcal{B}_j[t_{j_1}, t_{j_n}] \tag{4.3}$$

at 'levels' i, j (possibly $j = i$), each respecting the 'parts' to 'wholes' condition of (4.2). Both classifications lead to their corresponding diagrams as in Fig. 5c, as explained in §3.2, with an induced (logic) infomorphism $\mathbf{C}_i \longrightarrow \mathbf{C}_j$ between the cocone vertex classifications of the corresponding CCCDs derived from Fig. 5c (again, the time intervals are understood). Following the formalism of §3.4, we thus obtain induced (logic) infomorphisms:

$$\begin{gathered} \mathbf{Dg}_i[t_{i_1}, t_{i_n}] \longrightarrow \mathbf{Dg}_j[t_{j_1}, t_{j_n}] \\ \mathbf{C}_i \mapsto \mathbf{C}_j \end{gathered} \tag{4.4}$$

Schematically, this leads to a typical, i.e. generic, quasi-hierarchial configuration as depicted in Fig. 7 below. Note that the assumption of logic infomorphisms provides an underlying semantic structure to the various mechanisms as discussed in §2 and §3.

$$\begin{tikzcd}
\cdots & \cdots & \cdots & \cdots & \\
{} & \mathbf{Dg}_j[t_{j_1}, t_{j_n}] \arrow[dl, dotted] \arrow[d, dotted] & {} & {} \arrow[d, dotted] & \\
\mathbf{Dg}_w[t_{w_1}, t_{w_n}] \arrow[u, dotted] & \mathbf{Dg}_k[t_{k_1}, t_{k_n}] \arrow[dl, dotted] \arrow[d, dotted] & \mathbf{Dg}_\ell[t_{\ell_1}, t_{\ell_n}] \arrow[ul, dotted] \arrow[u, dotted] & \mathbf{Dg}_p[t_{p_1}, t_{p_n}] \arrow[l, dotted] \arrow[d, dotted] \arrow[r, dotted] & {} \\
\mathbf{Dg}_v[t_{v_1}, t_{v_n}] \arrow[u, dotted] & {} & \mathbf{Dg}_m[t_{m_1}, t_{m_n}] \arrow[ul, dotted] \arrow[u, dotted] \arrow[dl, dotted] & \mathbf{Dg}_q[t_{q_1}, t_{q_n}] \arrow[l, dotted] \arrow[dl, dotted] \arrow[d, dotted] \arrow[r, dotted] & {} \\
{} \arrow[u, dotted] & \mathbf{Dg}_s[t_{s_1}, t_{s_n}] \arrow[ul, dotted] \arrow[dl, dotted] & \mathbf{Dg}_u[t_{u_1}, t_{u_n}] \arrow[l, dotted] \arrow[u, dotted] \arrow[dl, dotted] & \mathbf{Dg}_r[t_{r_1}, t_{r_n}] \arrow[l, dotted] \arrow[dl, dotted] \arrow[d, dotted] \arrow[r, dotted] & {} \\
\cdots & \cdots & \cdots & \cdots &
\end{tikzcd}$$

Fig. 7: A typical complex of interactive CCCDs corresponding to a mereological object-token hierarchy that is maintained over time. Note that by taking the Chu space and Classification simplicial nerve construction of Fields and Glazebrook (2018, §5.4 -§6,6), the above diagram admits an underlying simplicial complex to which a range of simplicial methods (e.g. homotopies) can be applied.

The configurations depicted in Fig. 7 are not strictly hierarchial, even though the corresponding colimits are iterated. Why such a configuration cannot be strictly hierarchial is clear: a part can be a part of many wholes, and even of wholes at different levels (for instance, an employee can be part of a division, but also part of a company). As will be seen later, it is in this respect that a mereotopological complex differs from a standard category-theoretic hierarchy in the sense of e.g. Baas, Ehresmann and Vanbremeersch (2004); **?**. In particular, because of the existence of co-planar complexes and bi-directional arrows, it is not always the case that that a relevant object of level $n+1$, say, is the colimit of at least one diagram at level n. However, somewhat in line with Baas, Ehresmann and Vanbremeersch (2004), we may also consider restrictions or extensions of such a diagram that correspond to the "point of view" of an "observer" either internal the system or external (e.g. the "environment" of the system), through which "selection" for coherence or some other functional criterion induces the further levels of structure.

Lateral connections at each level of Fig. 7 indicate, e.g. probable co-occurrence in a scene. Vertical arrows indicate mereological inclusions going upwards and top-down predictions from current "understanding" or "prior probabilities" going downwards. These do not necessarily select the same relations, as they often do not in real-life situations. Cocones exist in both directions in this mereological structure: linked (time dependent) colimit cocones as underlying a typical CCCD, and CCCDs as linked by a network of logic infomorphims $... \longrightarrow \mathbf{C}_i \longrightarrow \mathbf{C}_{i+1} \longrightarrow \mathbf{C}_{i+2} \longrightarrow ...$ between the cocone vertex classifications (see Fields and Glazebrook, 2018, §6.4, §8).

As a simplified example, consider the two-layer mereological network in Fig. 8. "Complexes" here could consist of EL types such as [cat], [dog], [chair], etc. Each of these has many tokens that are specific individual cats, dogs, tables, etc. The "part" level can include visually-identifiable, but non-essential features of these types, such as "has four legs", "has fur", "walks with a gait X", "is white with brown spots", etc. as well as essential features. There are also tokens at this "part" level, e.g. four particular legs, a particular pattern of white-with-brown spots, etc. Cats, dogs and chairs can all have four legs, but exemplars with fewer legs are also possible. Each "part" token is linked to only a single "complex" token, e.g. a particular leg is a part of a particular cat, dog or table, though as noted above this uniqueness would fail if more abstract complexes were included. "Complex" types, and hence tokens, are related by co-occurrence links pertaining to scenes (i.e. higher-level complexes); "part" types, and hence tokens, are related by co-occurence links pertaining to complexes.

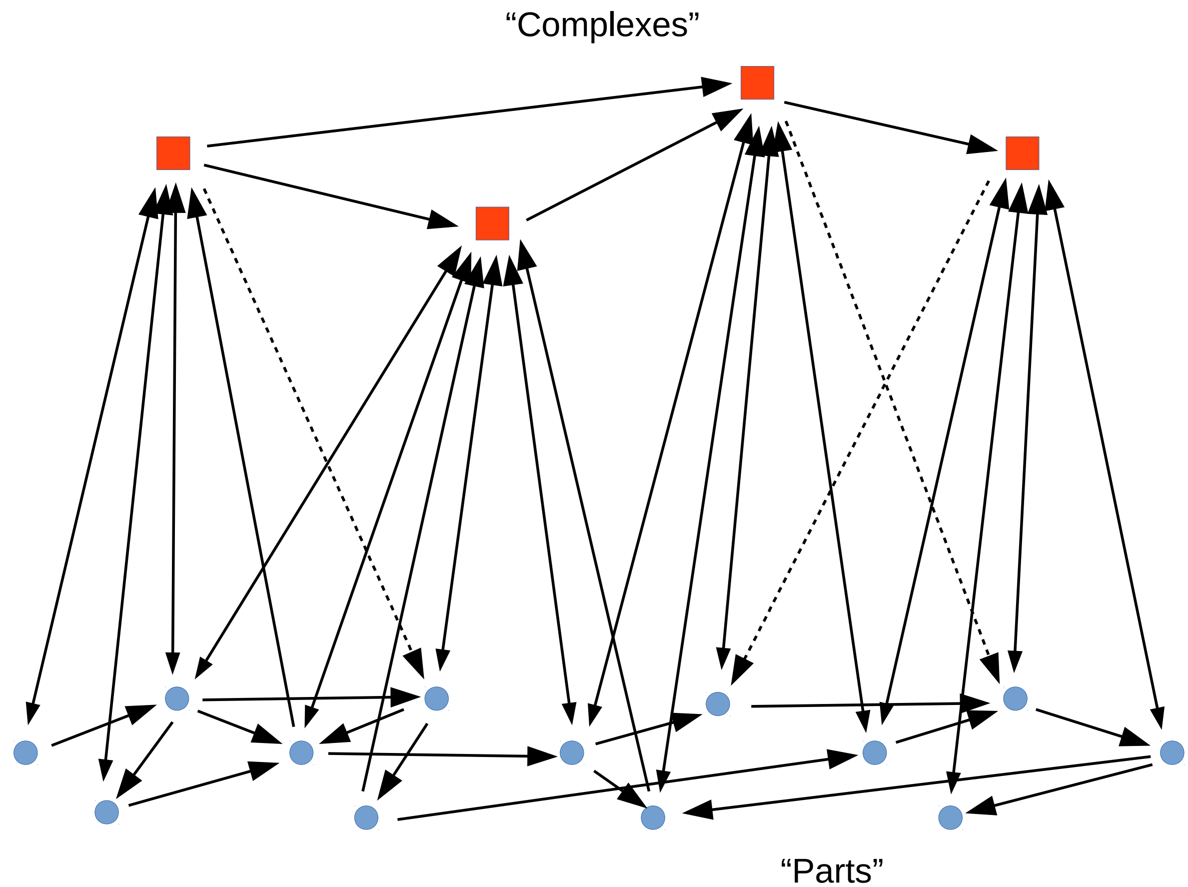

Fig. 8: A two-layer mereological network with "parts" at the lower level, and "complexes" at the upper level. Each level comprises both tokens and types. Mereological relations may be supported by only some exemplars (dashed lines).

"Learning" in the network of Fig. 8 would consist of: 1) associating altogether new lower- and upper-level types into the network; 2) distinguishing new high-level individuals as clusters of low-level tokens; 3) Convergence toward better predictions (i.e. all vertical arrows becoming bidirectional). "Binding" in this system is associating a collection of upward arrows with an upper-level token. "Abstraction" is the grouping of upper-level tokens into an upper-level type while preserving all arrows. Previous work demonstrating general models of ANNs (Fields and Glazebrook, 2018, §7.1) shows that such processes are allowed in principle; different specific choices of algorithms for these processes would be expected to produce different hierarchical structures. Closely related is the robotics example of Lyubova, Ivaldi and Filliat (2016) where a two-level representation is constructed along these lines from the bottom up, including the correct association of tokens at the "complex" level (representing a robot and an experimenter) with tokens at the "part" level (representing the robot's hand and the experimenter's hand).

4.3 From mereology to mereotopology: Distinguishing objects by boundaries

The detection of edges and their extension into contours that segment an image into bounded, non-overlapping regions is one of the earliest stages of visual processing (reviewed by Wagemans et al., 2012a). What, however, distinguishes a two- or even three-dimensional array of bounded, non-overlapping pattern elements – e.g. an array of colour or texture patches – from an array of bounded, non-overlapping objects? As discussed in §2.2 above, animacy, agency and independent manipulability are important indicators of bounded objecthood during infancy and early childhood when object categories are first being learned and populated with exemplars. What, however, are the inferences that enforce boundedness for objects, and how does the constraint of having a boundary affect the informational relations outlined above?

The key idea of mereotopology is that the parts of an object must be *inside* the object, i.e. contained within its boundary (Casati and Varzi, 1999; Smith, 1996). This constraint is, clearly, more easily satisfied for boundaries that are (at least approximately) smooth and convex. As simplicity and hence resource efficiency appear to be general principles of perceptual system organization (Wagemans et al., 2012b), one can expect perceivers to "see" smooth, convex boundaries – e.g. convex hulls of geometrically more complex objects – more easily. Imposing smoothness and convexity – e.g. by constructing the convex hull of a geometrically more complex object (e.g. as in Lyubova, Ivaldi and Filliat, 2016) – is a form of coarse-graining. We can, therefore, suggest that constructing an "exterior" boundary around a collection of parts that then serves as a boundary for the whole is a coarse-graining operation. Humans are, as noted in §2.2 above, highly accomplished at such boundary construction, especially for moving objects, with the ability to rapidly and accurately identify point-light walkers and similar disconnected displays as a compelling example. Static features are minimized by design in moving light displays such as point-light walkers in order to specifically probe dorsal visual stream processing. As the dorsal stream does not "see" shape (Flombaum, Scholl and Santos, 2008), the "human-shaped" boundary in this case is imposed from above, i.e. by the categorization process. Imposing this boundary coarse-grains the individually erratic, but highly correlated, motion of the individual point lights into bounded object motion from left to right or vice-versa. Inverting the display inhibits categorization and hence boundary imposition, and is standardly used as a negative control (Johnson and Hannon, 2015). If boundary construction is treated as coarse-graining, then the simplicial methods introduced in Fields and Glazebrook (2018, §5) are immediately applicable, and indeed provide a general method of constructing object boundaries from the bottom up in any mereological hierarchy representable in the CCCD form as in Fig. 7.

As noted earlier, a scene is a mereological complex; segmenting a scene by adding boundaries makes it a mereotopological complex. At the "top" of the mereotopological hierarchy, a whole scene can be considered a multilayer complex of simplicial complexes (i.e. identified "whole" objects) of simplices (identifiable "part" objects). Recall from Fields and Glazebrook (2018, §5.1) that any such complex, at any level of the hierarchy, has associated barycentric coordinates and a natural metric. Distances *within* a simplicial complex at level n of the hierarchy, however, can also be viewed as distances *between* its component simplices at level $n-1$. Boundaries, therefore, induce geometric relations between the bounded objects. In this sense, perceived spatial relations can be viewed as "emergent" from the distinctions between objects, a view with striking similarities to recent proposals within fundamental physics (cf. Fields et al., 2017).

4.4 Channels, inter-object boundaries and interactions

The spatial separation between objects induced by their boundaries – and hence by their distinguishability – generates a time-dependent exchange of information and hence an *interaction* as this term is traditionally understood. Again working from the top down in a mereotopological hierarchy of simplicial complexes indentified with local classifications, every channel between classifications at level n can also be viewed as a channel between the corresponding "objects", i.e. simplicial complexes. This channel corresponds to the boundary between the "objects" if they are adjacent, i.e. if the corresponding simplicial complexes share $(n-1)$-level faces. It is natural to think of the information transmitted along the channel as "encoded on" this boundary, i.e. as encoded holographically as this term is used in physics (Fields et al., 2017). If the objects are not adjacent, i.e. if the connecting channel is a composition at level n, the channel can be thought of as passing through a shared "environment" interposed between the objects. The components of the composed channel cannot in general be expected to be isomorphisms; hence the structure of this interposed environment affects the interaction between the objects.

Perceiving object motion requires tracking the identity of the "moving" object through time; hence it involves a temporal sequence of mereological hierarchies along the lines of Fig. 7. The structure of the top-level scene is different at each time increment; hence the metric relations between component simplicial complexes is time-dependent. Changing the relative positions of objects re-shapes their shared environment, in general altering the interaction between them. The distance and material, e.g. transparency or electrical permittivity, dependence of physical interactions can, therefore, be viewed as qualitatively "emergent" from the simplicial structure of classifications.

Coupling the perception of space and spatial relations, including distance-dependent interactions, to mereological reasoning in this way generates strong predictions, particularly about development and the consequences of damage. The representation of visual space and spatial relations is increasingly seen as a hippocampal function in mammals, including humans (see Moser, Kropff and Moser, 2008; Zeidman and Maguire, 2016, for reviews), coupling spatial representation anatomically to the construction of scenes/events and the encoding of episodic memories (cf. Fig. 1). While the architecture supporting this spatial representation is innate, the extent to which space must be learned through non-specific experience during infancy remains unknown. The current framework predicts that this hippocampal architecture also supports mereological reasoning, and that non-specific experience of mereological relationships is essential to a fully-developed representation of space. It predicts, in particular, that mereological relationships are not encoded solely by semantic memory systems, including the concept representation system implemented by ATP. If this is correct, developmental variants such as autism that affect spatial perception can be expected to also affect mereological reasoning; whether weak central coherence in autism (Happé and Frith, 2006) can be understood in this way remains to be determined. Hippocampal damage that affects spatial orientation and reasoning, e.g. consequent to Alzheimer's disease (Jicha and Carr, 2010), should similarly impact mereological reasoning. To our knowledge, correlational studies along these lines have not yet been performed.

4.5 Allocating attention to parts and wholes

Treating visual objects as multi-level mereological complexes consisting of parts and wholes raises immediate questions regarding attentional control. Top-down, gestalt-oriented perceptual processing can be expected to proceed from the top of the mereological hierarchy at which global structures are encoded toward the bottom of the hierarchy where local details about parts are encoded. Bottom-up, detail-oriented processing – the processing by which scenes are constructed from local features such as edges and colours – proceeds in the opposite direction. Entry-level objects, including both manipulable objects such as tools and social objects such as other people, lie somewhere in between. It is such entry-level objects that are most familiar, most goal-relevant, and most attractive of both dorsal and ventral attention. Attention can, therefore, be expected to be captured at the entry level, while delaying response to both global and detailed "part" information. This convergence of attention is, in humans, a compromise between the right and the left hemispheres, the left emphasizing information from higher spatial frequencies subserving an analysis of localized detail with lessened perception of salience, and the right emphasizing lower spatial frequencies subserving the more rapid global processing of information, heightened attention to salience, and eventually, the over-riding holistic organization of scenes and events (Bar, 2004; Kimchi, 2015). How this compromise plays out varies between individual perceivers and perceptual contexts, as the Navon task (Navon, 1977) and similar experiments demonstrate.

How human perception constructs the mereological complexes perceived as "whole" objects above the entry level, and hence how mereological categories above the entry level are learned, remains poorly understood. Such complexes lie, however, at the heart of social cognition, and their accurate perception and interpretation were presumably critical to human evolution (see Adolphs, 2009, for review). Social robotics clearly faces a similar problem in uncircumscribed material and social environments. Whether gestalt principles for the integration of local information can "emerge" heuristically from association measures such as probability or "simplicity" remains unclear (Wagemans et al., 2012b), as do criteria for the time persistence of mereological complexes. Local elements of a hierarchial pattern are not, moreover, in general properties of the global form; they are not necessarily parts of the whole in themselves. Thus a global processing advantage is not, *per se*, an advantage of a global property of a visual object/event over its local properties, but rather an attentional advantage of higher-level over lower-level processing. While autism provides cases in which a global processing advantage systematically fails or never develops, the presumably-existing functionally opposite population of individuals for whom perception is highly biased toward global processing does not have a syndromic name and has not been systematically investigated.

5 Conclusion

In Part I of this work (Fields and Glazebrook, 2018), we reviewed the rich set of tools that Chu spaces and Channel Theory provide for investigating relationships between informational structures and representing semantic information flows between such structures. These and other methods of category theory have been applied widely in computer science, and are seeing increasing applications in physics. As reviewed in Fields and Glazebrook (2018), applications of category theoretic methods in the cognitive sciences – mainly in the investigation of ontologies and ontology convergence – have mainly been carried out at a high level of abstraction. In this Part II, we have begun the process of characterizing object perception in category-theoretic terms, particularly in terms of

Chu spaces, classifications, simplicial complexes, and local logics, at the level of detail allowed by empirical results of both cognitive and neurocognitive investigations. These formal methods provide a natural and intuitively-clear representation of object perception as a multi-stage process in which tokens at one level serve as types at a higher level. The cone – co-cone duality captured in CCCDs is particularly useful as a representation of the bidirectional information flow employed by the predictive coding systems that the brain appears to implement at multiple scales. Not only do tokens at each scale become types at the next, low-level tokens also serve as "types" that organize and impose consistency conditions on higher-level types, which serve as the "tokens" in this dual relationship. Types and tokens thus play dual roles at each scale. We conclude, therefore, that visual object perception can be considered scale-free in this sense, at least approximately and within limits set by attentional and inferential abilities, and suggest that such a scale-free organization may characterize cognition more generally. The analysis we present in §3 extends previous work on the representation of ANNs by making this scale-independent CCCD duality explicit. It also makes explicit the essential role of inferences of FCHs – here captured as infomorphisms – in tracking object identity through time.

Human beings, and presumably other animals with relatively complex cognitive systems, employ both abstraction and mereological hierarchies to categorize objects. We showed in §4 that networks of time-indexed CCCDs provide a natural representation of the mutually-constraining relationship between these two categorization methods, particularly as they are employed in object-identity tracking. We then explored briefly the emergence of spatial relationships and interactions between objects from their description as simplicial complexes embedded in the larger simplicial complex that constitutes a perceptual scene. This emergence-based approach to mereotopology differs significantly from previous approaches that are geared toward *a priori* specification of ontologies (Lê and Janicki, 2008; Smith, 1996). It is worth noting that our categorical analysis of mereological reasoning, or of perception in general, makes no claims about the nature of consciousness and does not commit the "mereological fallacy" of attributing human-scale psychological properties to neurocognitive components (Bennett and Hacker, 2003). Indeed it shows why such a fallacy is a fallacy: what is scale-free in a CCCD is not any particular attribute, but the multiple roles played by the satisfaction relation $\Vdash$ linking within-scale types and tokens.

This initial foray into the categorical representation of cognitive processes at depth raises a number of questions and illuminates several open problems. One of the deepest is whether the satisfaction relations $\Vdash$ operating between tokens and types at any of the processing levels considered here are well-defined. While it must be assumed that they can be defined formally to develop models, it remains possible that "$\Vdash$" is token, type, context, or time dependent, as studies of the dependence of language on unspecified "background knowledge" (Searle, 1983) or of cognition generally on "embodiment" (Anderson, 2003; Chemero, 2013) might suggest. If this is the case, there *are no* formally-specifiable criteria for either object persistence or mereological composition, and any proposed criteria must be viewed as purely heuristic. Problems that require further work include: 1) the implications of the present results for the representation of ontologies and particularly for ontology convergence between agents that have encountered non- or only partially-overlapping collections of individual objects, 2) the extent to which local logics define or constrain semantic relations between either tokens or types, 3) the extent to which cognitively-significant differences in spatial scale can be captured by coarse-graining, and 4) the question of why the geometry emergent from human visual perception should be three-dimensional. There also remains the open question of how cognition is affected when one or the other of the categorization systems breaks down, as appears to be the case with the mereological system in autism where a sense of context, and an overall gestalt in certain situations may be adversely impacted (Happé and Frith, 2006).

Acknowledgement

The authors wish to thank two anonymous referees for constructive criticism and helpful comments which assisted in improving the overall presentation

• The authors report no conflict of interest involved in this work.

References

Adolphs, R. (2009). The social brain: Neural basis for social knowledge. *Annual Review of Psychology, 60*, 693–716.

Alain, C., Arnott, S. R., Hevenor, S., Graham, S. & Grady, C. L. (2001). "What" and "where" in the human auditory system. *Proc. Natl. Acad. Sci. USA, 98*, 12301–12306.

Anderson, M. L. (2003). Embodied cognition: A field guide. *Artif. Intell., 149*, 91–130.

Baars, B. J & Franklin, S. (2003). How conscious experience and working memory interact. *Trends Cogn. Sci., 7*, 166–172.

Baars, B. J., Franklin, S. & Ramsoy, T. Z. (2013). Global workspace dynamics: Cortical "binding and propagation" enables conscious contents. *Front. Psychol., 4*, 200.

Baas, N. A., Ehresmann, A. C. & Vanbremeersch, J.-P. (2004). Hyperstructures and memory evolutive systems. *Int. J. Gen. Syst., 33*, 553–568.

Baillargeon, R. (2008). Innate ideas revisited: For a principle of persistence in infants physical reasoning. *Perspect. Psychol. Sci., 3*, 2–13.

Baillargeon, R., Stavans, M., Wu, D., Gertner, Y., Setoh, P., Kittredge, A. K. & Bernard, A. (2012). Object individuation and physical reasoning in infancy: An integrative account. *Lang. Learn. Devel., 8*, 4–46.

Bar, M. (2004). Visual objects in context. *Nat. Rev. Neurosci., 5*, 617–629.

Barr, M. (1979). *-Autonomous categories, with an appendix by Po Hsiang Chu. *Lecture Notes in Mathematics, 752.* Berlin: Springer.

Barr, M. (1991). *-Autonomous categories and linear logic. *Mathematical Structures in Computer Science, 1*, 159–178.

Barwise, J. & Seligman, J. (1997). *Information Flow: The Logic of Distributed Systems.* Cambridge Tracts in Theoretical Computer Science, 44. Cambridge, UK: Cambridge University Press.

Barwise, J. & Perry, J. (1983). *Situations and Attitudes.* Cambridge, MA: Bradford Books, MIT Press.

Bassett, D. S. & Bullmore, E. (2006). Small world brain networks. *Neuroscience, 12*, 512–523.

Bastos, A. M., Usrey, W. M., Adams, R. A.,Mangun, G. R., Fries, P. & Friston, K. J. (2012). Canonical microcircuits for predictive coding. *Neuron, 76*, 695–711.

Bauer, P. B. (2006). Constructing a past in infancy: A neuro-developmental account. *Trends Cogn. Sci., 10*, 175–181.

Bennett, M. & Hacker, P. (2003). *Philosophical Foundations of Neuroscience*. London: Blackwell.

Blake, R. & Shiffrar, M. (2007). Perception of human motion. *Ann. Rev. Psychol., 58*, 47–73.

Booth, R. D. L. & Happé, F. (2018). Evidence of reduced global processing in autism spectrum disorder. *J. Autism Devel. Disord., 48*, 1397–1408.

Brady, T. F., Konkle, T., Alvarez, G. A. & Oliva, A. (2008). Visual long-term memory has a massive storage capacity for object details. *Proc. Natl. Acad. Sci. USA, 105*, 14325–14329.

Casati, R. & Varzi, A. C. (1999). *Parts and Places: The Structures of Spatial Representations*. Cambridge, MA: MIT Press.

Chemero, A. (2013). Radical embodied cognitive science. *Rev. Gen. Psych., 17*, 145–150.

Churchland, P. M. (2012). *Plato's Camera: How the Physical Brain Captures a Landscape of Abstract Universals*. Cambridge, MA: MIT Press.

Clarke, A. & Tyler, L. K. (2015). Understanding what we see: How we derive meaning from vision. *Trends Cogn. Sci., 19*, 677–687.

Cloutman, L. L. (2013). Interaction between dorsal and ventral processing streams: Where, when and how? *Brain Lang., 127*, 251–263.

Dehaene, S. & Naccache, L. (2001). Towards a cognitive neuroscience of consciousness: Basic evidence and a workspace framework. *Cognition, 79*, 1–37.

Dehaene, S. & Changeux, J. P. (2004). Neural mechanisms for access to consciousness. In: Gazzaniga, M. S. (Ed.) *The Cognitive Neurosciences, 3rd Edn.* Cambridge, MA: MIT Press (pp. 1145–1157).

Dehaene, S., Charles, L., King, J.-R. & Marti, S. (2014). Toward a computational theory of conscious processing. *Curr. Opin. Neurobiol., 25*, 76–84.

Dretske, F. (1981). *Knowledge and the Flow of Information*. Cambridge MA: MIT Press.

Dretske, F. (2000). *Perception, Knowledge, and Belief: Selected Essays*. Cambridge Studies in Philosophy. New York: Cambridge University Press.

Eichenbaum, H., Yonelinas, A. R. & Ranganath, C. (2007). The medial temporal lobe and recognition memory. *Ann. Rev. Neurosci., 30*, 123–152.

Fei-Fei, L., Fergus, R. & Perona, P. (2006). One-shot learning of object categories. *IEEE Trans. Pattern Anal. Mach. Intell., 28*, 594–611.

Fields, C. & Glazebrook, J. F. (2018). A mosaic of Chu spaces and Channel Theory I: Category-theoretic concepts and tools. *J. Expt. Theor. Artif. Intell.* in press.

Fields, C. (2012). Trajectory recognition as the basis for object individuation: A functional model of object file instantiation and object token encoding. *Front. Psychol.: Percept. Sci., 2*, 49.

Fields, C. (2012). The very same thing: extending the object token concept to incorporate causal constraints on individual identity. *Adv. Cogn. Psychol., 8*, 234–247.

Fields, C. (2013). How humans solve the frame problem. *J. Expt. Theor. Artif. Intell., 25*, 441–456.

Fields, C. (2014). Motion, identity and the bias toward agency. *Front. Hum. Neurosci., 8*, 597.

Fields, C. (2016). Visual re-identification of individual objects: A core problem for organisms and AI. *Cogn. Process., 17*, 1–13.

Fields, C. (2017). Object Permanence. In: Shackelford, T. K. & Weekes-Shackelford, V. A. (Eds.), *Encyclopedia of Evolutionary Psychological Science.* New York: Springer, (Chap. 2373).

Fields, C., Hoffman, D. D., Prakash, C. & Prentner, R. (2017). Eigenforms, interfaces, and holographic encoding: Toward an evolutionary account of objects and spacetime. *Construct. Found., 12*, 265–273.

Fields, C., Hoffman, D. D., Prakash, C. & Singh, M. (2018). Conscious agent networks: Formal analysis and application to cognition. *Cogn. Syst. Res., 47*, 186–213.

Fievet, L. & Sornette, D. (2018). Calibrating emergent phenomena in stock markets with agent based models. *PLoS One, 13*, e0193290.

Firestone, C. & Scholl, B. J. (2016). Cognition does not affect perception: Evaluating the evidence for "top-down" effects. *Behav. Brain Sci., 39*, e229.

Flombaum, J. I., Scholl, B. J. & Santos, L. R. (2008). Spatiotemporal priority as a fundamental principle of object persistence. In Hood, B. & Santos, L. (Eds.), *The Origins of Object Knowledge.* New York: Oxford University Press (pp. 135–164).

Franklin, S. & Patterson, F. G. J. Jr. (2006). The LIDA architecture: adding new modes of learning to an intelligent, autonomous, software agent. *IDPT-2006 Proceedings (Integrated Design and Process Technology).* Society for Design and Process Science (pp. 1–8).

Franklin, S., Madl, T., D′Mello, S. & Snaider, J. (2014). LIDA: a systems-level architecture for cognition, emotion and learning. *IEEE Trans. Auton. Mental Devel., 6*, 19–41.

Friston, K. J. (2010). The free-energy principle: A unified brain theory? *Nat. Rev. Neurosci., 11*, 127–138.

Friston, K. J. & Kiebel, S. (2009). Predictive coding under the free-energy principle. *Phil. Trans. R. Soc. London B, 364*, 1211–1221.

Friston, K. J., Kilner, J. & Harrison, L. (2006). A free energy principle for the brain. *J. Physiology – Paris, 100*, 70–87.

Friston, K. J., Rigoli, F., Ognibene, D., Mathys, C., FitzGerald, T. & Pezzulo, G. (2015). Active inference and epistemic value. *Cogn. Neurosci., 6*, 187–214.

Gao, W., Alcauter, S., Smith, J. K., Gilmore, J. H. & Lin, W. (2015). Development of human brain cortical network architecture during infancy. *Brain Struct. Funct., 220*, 1173–1186.

Glazebrook, J. F. & Wallace, R. (2009). Small worlds and red queens in the global workspace: An information-theoretic approach. *Cogn. Syst. Res., 10*, 333–365.

Goodale, M. A. (2014). How (and why) the visual control of action differs from visual perception. *Proc. R. Soc. B, 281*, 20140337.

Goodale, M. A. & Milner, A. D. (1992). Seperate visual pathways for perception and action. *Trends Neurosci., 15*, 20–25.

Grossberg, S. (1980). How does a brain build a cognitive code? *Psychol. Rev., 87*, 1–51.

Grossberg, S. (1988). Nonlinear neural networks: Principles, mechanisms, and architectures. *Neural Networks, 1*, 17–61.

Grossberg, S. (2007). Towards a unified theory of neocortex: Laminar cortical circuits for vision and cognition. *Prog. Brain Res., 165*, 79–104.

Grossberg, S. (2013). Adaptive Resonance Theory: How a brain learns to consciously attend, learn, and recognize a changing world. *Neural Networks, 37*, 1–47.

Hampton, J. A. (2016). Categories, prototypes and exemplars. In: N. Reimer (Ed.), *Routledge Handbook of Semantics*. Abingdon, UK: Routledge (Ch. 7).

Happé, F. & Frith, U. (2006). The weak coherence account: Detail-focused cogitive style in autism spectrum disorders. *J. Autism Devel. Disord., 36*, 5–25.

Hayne, H. (2004). Infant memory development: Implications for childhood amnesia. *Devel. Rev., 24*, 33–73.

Hochstein, S. & Ahissar, M. (2002). View from the top: hierarchies and reverse hierarchies in the visual system. *Neuron, 36*, 791–804.

Hoffman, D. D., Singh, M. & Prakash, C. (2015). The interface theory of perception. *Psychonom. Bull. Rev., 22*, 1480–1506.

Hommel, B. (2004). Event files: feature binding in and across perception and action. *Trends Cogn. Scii.,8*, 494–500.

Huang, H., Shu, N., Mishra, V., Jeon, T., Chalak, L., Wang, Z. J., Rollins, N., Gong, G., Cheng, H., Peng, Y., Dong, Q. & He, Y. (2015). Development of human brain structural networks through infancy and childhood. *Cereb. Cortex, 25*, 1389–1404.

Husserl, E. (1970). *Logical Investigations*. London: Routledge and Keagan Paul.

Jicha, G. A. & Carr, S. A. (2010). Conceptual evolution in Alzheimer's disease: Implications for understanding the clinical phenotype of progressive neurodegenerative disease. *J. Alzheimer's Dis., 19*, 253–272.

Johnson, S. P. & Hannon, E. E. (2015). Perceptual development. In Liben, L. S. & Müller, U. (Eds.) *Handbook of Child Psychology and Developmental Science, Vol. 2: Cognitive Processes* Hoboken, NJ: Wiley (pp. 63–112).

Kahneman, D. (2011). *Thinking, Fast and Slow.* Penguin, London.

Kahneman, D., Triesman, A. & Gibbs, B. J. (1992). The reviewing of object files: Object-specific integration of information. *Cogn. Psychol., 24*, 175–219.

Kalfoglou, Y. & Schorlemmer, M. (2003). IF-Map: An ontology-mapping method based on information-flow theory. *J. Data Semantics I* (Lecture Notes in Computer Science). Berlin: Springer (pp. 107–127).

Keifer, M. & Pulvermüller, F. (2012). Conceptual representations in mind and brain: Theoretical developments, current evidence and future directions. *Cortex, 7*, 805–825.

Kent, R. E. (2016). Information flow in logical environments. Preprint arXiv:1603.03475v1[cs.LO].

Kimchi, R. (2015). The perception of hierarchial structure. In: Wagemans, J. (Ed.), *Oxford Handbook of Perceptual Organization.* Oxford, UK: Oxford University Press (pp. 129–149).

Kotseruba, I. & Tsotsos, J. K. (2018). A review of 40 years in cognitive architecture research: Core cognitive abilities and practical applications. Preprint arxiv:1610:08602v3 [cs.AI].

Kristan, M., Matas, J., Leonardis, A. ... Pflugfelder, R. (2015). The visual object tracking VOT2015 challenge results. *IEEE Int. Conf. Comp. Vision (ICCV)*, Workshops (pp. 1–23).

Lando, G. (2017). *Mereology: A Philosophical Introduction.* London: Bloomsbury.

Lawson, R. P., Rees, G. & Friston, K. J. (2014). An aberrant precision account of autism. *Front. Human Neurosci., 8*, 302.

Lê, D. T. M. & Janicki, R. (2008). A categorical approach to mereology and its applications to modeling software components. In: Peters, J. F. & Skowron, A. (Eds.), *Transactions on Rough Sets VIII.* (Lecture Notes in Computer Science, 5084). Berlin: Springer (pp. 146–174).

Leśniewski, S. (1929). Grundzüge eines neuen Systems der Grundlagen der Mathematik. *Fundamenta Mathematicae, 24*, 1–81.

Lyubova, N., Ivaldi, S & Filliat, D. (2016). From passive to interactive object learning and recognition through self-identification on a humanoid robot. *Auton. Robots, 40*, 33–57.

Mahon, B. Z & Caramazza,A. (2009). Concepts and categories: A cognitive neuropsychological perspective. *Ann. Rev. Psychol., 60*, 27–51.

Maloney, L. T. & Zhang, H. (2010). Decision-theoretic models of visual perception and action. *Vision Res., 50*, 2362–2374.

Marr, D. (1982). *Vision: A Computational Investigation into the Human Representation and Processing of Visual Information.* San Francisco, CA: W. H. Freeman.

Martin, A. (2007). The representation of object concepts in the brain. *Ann. Rev. Psychol., 58*, 25–45.

Moser, E. I., Kropff, E. & Moser, M.-B. (2008). Place cells, grid cells, and the brains spatial representation system. *Ann. Rev. Neurosci., 31*, 69–89.

Muehlhaus, J., Heim, S.,Altenbach, F., Chatterjee, A., Habel, U. & Sass, K. (2014). Deeper insights into semantic relations: An fMRI study of part-whole and functional associations. *Brain & Lang., 129*, 30–42.

Navon, D. (1977). Forest before trees: The precedence of global features in visual perception. *Cogn. Psychol., 9*, 353–383.

Nguyen, S. M., Ivaldi, M., Lyubova, N., Droniou, A., Gérardeaux-Viret, D. Filliat, D., Padois, V., Sigaud, O. & and Oudeyer, P.-Y. (2013). Learning to recognize objects through curiosity-driven manipulation with the iCub humanoid robot. *Proceedings of the 2013 IEEE Third Joint International Conference on Development and Learning and Epigenetic Robotics (ICDL)* (pp. 1–8).

Nichols, S. & Bruno, N. (2010). Intuitions about personal identity: An empirical study. *Philos. Psychol., 23*, 293–312.

Oudeyer, P.-Y., Baranes, A. & Kaplan, F. (2013). Intrinsically motivated learning of real world sensorimotor skills with developmental constraints. In: Baldassarre, G. & Mirolli, M. (Eds.) *Intrinsically Motivated Learning in Natural and Artificial Systems.* Berlin: Springer (pp. 303–365).

Patel, H., Kothari, A. & Bhurchandi, K. (2015). 3-D face recognition: Features, databases, algorithms and challenges. *Artif. Intell. Rev., 44*, 393–441.

Pratt, V. (1995). Chu spaces and their interpretation as concurrent objects. *Lect. Notes in Comp. Sci., 1000* (1995), 392–405.

Pratt, V. (1999). Chu spaces. *School on Category Theory and Applications (Coimbra 1999)*, Vol. 21 of *Textos Mat. Sér. B*, University of Coimbra, Coimbra. (pp. 39–100).

Pratt, V. (1999). Chu spaces from the representational viewpoint. *Ann. Pure Appl. Logic, 96*, 319–333.

Pylyshyn, Z. (1984). *Computation and Cognition.* Cambridge, MA: MIT Press.

Rakison, D. H. & Yermolayeva, Y. (2010). Infant categorization. *WIRES Cogn. Sci., 1*, 894–905.

Ranganath, C. (2010). A unified framework for the functional organization of the medial temporal lobes and the phenomenology of episodic memory. *Hippocampus, 20*, 1263–1290.

Renoult, L., Davidson, P. S. R.,Palombo, D. J., Moscovitch, M. & Levine, B. (2012). Personal semantics: At the crossroads of semantic and episodic memory. *Trends Cogn. Sci., 16*, 550–558.

Ritchey, M., Libby, L. A. & Ranganath, C. (2015). Cortico-hippocampal systems involved in memory and cognition: The PMAT framework. *Prog. Brain Res., 219*, 45–64.

Rubinov, M. & Sporns, O. (2010). Complex network measures of brain connectivity: Uses and interpretations. *NeuroImage, 52*, 1059–1069.

Rugg, M. D. & Vilberg, K. L. (2013). Brain networks underlying episodic memory retrieval. *Curr. Opin. Neurobiol., 23*, 255–260.

Sathian, K., Lacey, S., Stilla, R.,Gibson, G. O., Deshpanda, G., Hu, X., LaConte, S. & Glielmi, C. (2011). Dual pathways for haptic and visual perception of spatial and texture information. *Neuroimage, 57*, 462–475.

Schaller, S. (2012). *A Man Without Words*, 2nd Ed. Berkeley, CA: University of California Press.

Schlesinger, M., Amso, D.,Johnson, S. P., Hantehzadeh, N. & Gupta, L. (2012). Using the iCub simulator to study perceptual development: A case study. In: *Proceedings of the 2012 IEEE international Conference on Development and Learning and Epigenetic Robotics (ICDL)* (pp. 1–6).

Scholl, B. J. (2007). Object persistence in philosophy and psychology. *Mind & Lang., 22*, 563–591.

Schwabe, L., Nader, K. & Pruessner, J. C. (2014). Reconsolidation of human memory: Brain mechanisms and clinical relevance. *Biol. Psychiatry, 76*, 274–280.

Searle, J. R. (1983). *Intentionality.* Cambridge, UK: Cambridge University Press.

Seligman, J. (2009). Channels: From logic to probability. In: Sommaruga, G. (Ed.) *Formal Theories of Information.* (Lecture Notes in Computer Science, 5363) Berlin: Springer (pp. 193–233).

Simons, P. (1987). *Parts: A Study in Ontology.* Oxford, UK: Oxford University Press.

Smith, B. (1996). Mereotopology: A theory of parts and boundaries. *Data Knowl. Engin., 20*, 287–303.

Sowa, J. F. (2006). Semantic networks. In: Nadel, L. (Ed.) *Encyclopedia of Cognitive Science.* New York: Wiley.

Spivak, D. I. & Kent, R. E. (2012). Ologs: A categorical framework for knowledge representation. *PLoS ONE, 7*, e24274.

Sporns, O. (2013). Network attributes for segregation and integration in the human brain. *Curr. Opin. Neurobiol., 23*, 162–171.

Sporns, O & Honey, C. J. (2006). Small worlds inside big brains. *Proc. Natl. Acad. Sci. USA, 111*, 15220–15225.

Spratling, M. W. (2008). Predictive coding as a model of biased competition in visual selective attention. *Vision Res., 48*, 1391–408.

Spratling, M. W. (2016). Predictive coding as a model of cognition. *Cogn. Proc., 17*, 279–305.

Spratling, M. W. (2017). A hierarchical predictive coding model of object recognition in natural images. *Cogn. Comput., 9*, 151-167.

Treisman, A. (2006). Object tokens, binding and visual memory. In Zimmer, H. D., Mecklinger, A. & Lindenberger, U. (Eds.) *Handbook of Binding and Memory: Perspectives from Cognitive Neuroscience*. Oxford, UK: Oxford University Press (pp. 315–338).

Uddin, L. Q. (2015). Salience processing and insular cortical function and dysfunction. *Nat. Rev. Neurosci., 16*, 55–61.

Van de Cruys, S., Evers, K., Van der Hallen, R. ... Wagemans, J. (2014). Precise minds in uncertain worlds: Predictive coding in autism. *Psychol. Rev., 121*, 649–675.

Van Essen, D. C., Anderson, C. H. & Felleman, D. J. (1992). Information processing in the primate visual system: An integrated systems approach *Science, 255*, 419–423.

Vossel, S., Geng, J. J. & Fink, G. R. (2014). Dorsal and ventral attention systems: Distinct neural circuits but collaborative roles. *Neuroscientist, 20*, 150–159.

Wagemans, J., Elder, J. H., Kubovy, M., Palmer, S. E., Peterson, M. A., Singh, M. & von der Heydt, R. (2012). A century of Gestalt psychology in visual perception: I. Perceptual grouping and figureground organization. *Psych. Bull., 138*, 1172– 1217.

Wagemans, J., Feldman, J., Gepshtein, S., Kimchi, R., Pomerantz, J. R., van der Helm, P. A. & van Leeuwen, C. (2012). A century of Gestalt psychology in visual perception: II. Conceptual and theoretical foundations. *Psych. Bull., 138*, 1218– 1252.

Winkielman, P., Halberstadt, J., Fazendeiro, T. & Catty, S. (2006). Prototypes are attractive because they are easy on the mind. *Psychol. Sci., 17*, 799–806.

Zeidman, P. & Maguire, E. A. (2016). Anterior hippocampus: The anatomy of perception, imagination and episodic memory. *Nat. Rev. Neurosci., 17*, 173–182.

Zimmer, H. D. & Ecker, U. K. D. (2010). Remembering perceptual features unequally bound in object and episodic tokens: Neural mechanisms and their electrophysiological correlates. *Neurosci. Biobehav. Rev., 34*, 1066–1079.