Journal of Experimental and Theoretical Artificial Intelligence
In press, August 2012

How humans solve the frame problem

Chris Fields

815 East Palace Ave. # 14, Santa Fe, NM 87501 USA
fieldsres@gmail.com

Abstract: Both standard formulations of the frame problem and standard solutions implicitly assume that the re-identification of objects as persisting individuals between pre- and post-action contexts is unproblematic. In the case of human beings, this assumption is false: humans dedicate considerable cognitive resources to object re-identification. An analysis of both the phenomenology and neurocognitive implementation of object re-identification is used to show that in humans, all of the information architecturally available to solve the frame problem is in fact deployed for object re-identification. The frame problem is, therefore, equivalent to the object re-identification problem in the case of human problem solving.

Keywords: Cognition; Perception; Categorization; Object identification; Causal reasoning; Memory

Introduction

From its introduction by McCarthy and Hayes (1969) through the publication of *The Robot's Dilemma* (Pylyshyn, 1987) and its successor *The Robot's Dilemma Revisited* (Ford & Pylyshyn, 1996), the frame problem occupied center stage as perhaps *the* foundational problem of artificial intelligence. All of the major conceptual issues, from the nature of semantics to the tractability of planning, seemed to be tied up with the seemingly-simple problem of how to walk across the room without worrying about whether doing so would change the color of the walls or the air temperature outside. With the dawn of the 21st century, however, active discussion of the frame problem began to wane. Positions perhaps hardened as the AI and cognitive science communities fragmented into more clearly-defined schools of thought, with evolutionary psychologists siding with traditional representationalists in viewing heuristic satisficing as good enough for all practical purposes, dynamical systems and embodied cognition theorists tending to side with ecological realists and phenomenologists in claiming that all representational solutions suffer from infinite regress, and thorough-going rationalists claiming both that the frame problem is computationally intractable and that its solution by humans proves that human minds are not computational entities. Recent philosophical discussions tend to defend long-established positions with only minor elaborations (e.g. Fodor, 2008; Wheeler, 2008; Samuels, 2010). A recent review (Shanahan, 2009) succinctly summarizes the conceptual standoff.

What has been largely missing from discussions of the frame problem, even as it affects human problem solving, is explicit consideration of neuroscience. An exception to this is the work of Shanahan and Baars (2005), who address the question of how the human brain actually solves – to whatever extent it does actually solve – the frame problem in ordinary, resource-limited task environments. Shanahan and Baars address the rationalist worry that belief updating is intrinsically isotropic – that any belief can in principle influence the updating of any other belief – by arguing that at the level of description relevant to conscious thinking, cognition *is* isotropic, and that "relevance" is determined solely by the amount of attentional amplification a given representation can accumulate within the time alloted for reaching a solution. The present paper poses a complementary question: what information about a given belief is considered relevant by default, and how is this information represented? It argues that for humans, the information that is considered relevant by default in any problem solving situation involving actions carried out over time is the information required to solve a particular instance of the frame problem, the problem of determining whether the *identities* of the objects present in a remembered context are preserved in a newly-encountered context. Cognitive and neuroimaging studies focused largely on human visual object categorization and identification support a model in which the pre-motor system constructs unobserved and hence fictive causal histories that connect retrieved episodic memories of previous events to current observations (Fields, 2012). Such fictive causal histories (FCHs) must be constructed for every object that is re-identified as the same individual that participated in a remembered event. As discussed below, the relative contributions of feature and motion information to FCH construction are sensitively dependent on the time between observations, the category of object, and the observational context; "Leibniz's Law" that indiscernible objects are identical cannot be treated as the human default for object re-identification, even in the limit of very brief gaps in observation. If the FCH construction model is correct, all of the information that is architecturally available to solve the frame problem must be deployed, in humans, to solve the object re-identification problem; hence these two problems are equivalent for human problem solvers. This result suggests that the development of AI systems that solve the frame problem in a human-like way will require the replacement of the ubiquitous implicit assumptions that objects persist through time and are unproblematically re-identifiable using time-stable features with an explicit process that determines what objects in any given context are the same individual things that have been encountered previously. Viewed more broadly, it suggests that the assumption that objects are unproblematically re-identifiable is a gross oversimplification under any circumstances.

The paper is organized as follows. The next section, "The frame problem and the problem of object re-identification" distinguishes the traditional formulation of the frame problem in terms of property updating from a more radical version that involves the explicit re-identification of objects as persisting individual things from one context to the next. The third section, "How humans visually re-identify objects: Phenomenology" reviews the human use of featural and causal information in both short-term, within-scene visual object tracking and longer-term, between-scene object re-identification. The discussion focuses on visual object re-identification because vision is by far the best studied sensory modality; there are, however, no reasons to suggest that other senses are intrinsically more reliable that vision as indicators of object identity. The fourth section, "How humans visually re-identify objects: Implementation" reviews neurofunctional studies that relate object representations at increasing levels of integration to activity patterns in distributed networks of anatomically-characterized structures. It describes how humans deploy categorical constraints to construct FCHs that enable object re-identification across temporally and hence causally separated contexts. The fifth section, “The human solution: Default causal knowledge and resource-limited isotropy” shows that humans employ all of the information that is architecturally available to solve the frame problem in the FCH construction required to imagine possible outcomes. The paper concludes by suggesting that it is the fundamentally analogical character of pre-motor problem solving that renders the human solution of the frame problem adequate in practical contexts.

The frame problem and the problem of object re-identification

Suppose you must fly from the U.S. to Europe for a week-long meeting. After a harried morning at work, you have a quick lunch with your spouse and get on the plane. You arrive in Frankfurt, rumpled and jittery, the next morning. Just outside Immigration, you spot someone walking toward you whose features appear to perfectly match those of your spouse. You are filled immediately (within 400 ms; Eichenbaum, Yonelinas & Ranganath, 2007) with a feeling of familiarity. You are now faced with an instance of the frame problem, and your first thought, “Wow, that person looks just like my wife/husband” or “How did she/he get here?” indicates how you solved it. Your solution has systemic consequences. If you identified the person as your spouse, you are likely to experience confused delight when they rush to give you a hug; if you did not, you are likely to experience something akin to terror.

The frame problem is traditionally presented as the problem of circumscribing the set of objects for which properties need to be updated following an action. In the example above, one must decide whether the location of one's spouse needs to be updated following a flight to Europe. This traditional presentation rests on an implicit assumption: the assumption that objects can be unproblematically identified as the same or different between the pre-action and post-action contexts. If object identification is unproblematic, the frame problem is a problem about properties and the conditions under which they can be expected to change. As the example shows, however, this assumption of unproblematic object identification can fail: what can be uncertain – radically so – is whether an object with familiar properties that is encountered after an action is the *same thing* that had those properties before the action. If the assumption of unproblematic object identification is rejected, the frame problem becomes far more serious, as it now includes the problem of circumscribing the set of objects that must be explicitly re-identified as the same individual thing in the new context that follows an action. To update a famous example of Fodor's (from Pylyshyn, 1987), it becomes the problem of determining what components of the previous universe, if any, you are likely to encounter after you've turned off your fridge.

The implicit human assumptions that objects persist through time and can be unproblematically re-identified as the very same individual from one experienced context to the next is so basic as to be considered innate (Baillargeon, 2008); re-identifying objects is commonly treated a “primitive and nonconceptual” (Pylyshyn, 2009; p. 13) function that the early visual system is “‘wired’ to do” (p. 32). Most AI programs that deal with external objects are “wired” to regard those objects as persistent and re-identifiable through the use of fixed designators – effectively, proper names – that are regarded, within the chosen programming language, as having fixed semantics; the ubiquitous use of bar code labels or other easily-inspected unique-by-design features to encode fixed, unique designators exemplifies this semantic practice. If two such “rigid” designators are associated with the same identification criterion – for example, the same bar code – they are fully interchangeable as required by Leibniz's Law. With these semantic assumptions, the problem of deciding which of two identically-featured objects one has interacted with previously cannot even be coherently formulated: objects that satisfy the implemented semantic criteria are re-identified, while objects that do not satisfy them are not. In the case of humans, however, the semantic relationship between any given bit of neural "wiring" and the external world is not obvious, and may not even be discoverable through non-destructive testing of internal components. The human implementation of object re-identification must, therefore, be investigated at the level of overt behavior, including object-directed actions and verbal reports. Such experimental investigation of human object re-identification criteria has been on-going at least since the work of Heraclitus (c. 500 BCE); what has been learned is that human object re-identification is a complex, resource-intensive cognitive process.

How humans visually re-identify objects: Phenomenology

The modern era of object re-identification studies began with the work of Burke (1952), who showed that whether an object that briefly disappears from sight while passing through a tunnel is re-identified as "the same thing" after emerging depends on how long it spends in the tunnel. Ever more sophisticated experiments over the subsequent decades have established a number of robustly-supported conclusions regarding human within-scene visual object re-identification (reviewed by Treisman, 2006; Flombaum, Scholl & Santos, 2008; Gerhardstein *et al.*, 2009; Fields, 2011a). First, whether a briefly occluded object is re-identified as the same thing following the occlusion depends on the details of its speed and trajectory: some curvilinear trajectories at constant or variable speed indicate sameness, while others do not. Second, the number and kinds of trajectories that indicate sameness change from early infancy until about two years of age, after which they remain essentially fixed. Third and most significantly, the determination of object sameness is independent of changes in the object's static, motion-independent features – its size (as long as it remains finite), color, shape, etc. Trajectory information dominates static feature information in the "object file" (Treisman, 2006), the short-term memory (STM) resident representation on which downstream categorization and identification processes act. Using "objects" that were simply moving regions of random-dot displays, Gao and Scholl (2010) have shown that static features are in fact unnecessary for within-scene object re-identification; object files are constructed and correctly processed even for objects that cannot be distinguished from their backgrounds when not moving.

The frame problem is of minimal relevance during the brief periods of passive observation characterized by the experiments briefly reviewed above. It becomes relevant, however, as soon as object-focused actions and even few-second gaps in observation are introduced. In real-world tasks such as freeway driving, where the trajectories of surrounding objects cannot all be monitored simultaneously, quickly-distinguishable static features such as size and color become all-important in tracking object identities (Hollingworth and Franconeri, 2009). The successful driver's frame-problem solution never introduces the possibility that a lane change, for example, alters the color of a nearby car; visual detection of a car with an unexpected color is always treated as indicating the presence of a novel object that requires attention and possibly compensatory action. The relevance of features to object re-identification in temporally-extended interaction contexts such as driving indicate that a richer, more persistent representation than the object file is being acted upon. This more persistent representation has been labeled the "object token" by Zimmer and Ecker (2010). From a phenomenological perspective, object tokens differ from object files by incorporating category information that stabilizes features by organizing and naming them. Object tokens represent not moving blobs of color, but multiply-featured objects categorized as cars. Categorization reverses the priorities of motion and feature information from those in the object file; in object tokens, static features dominate motion. Unlike object files, object tokens may be maintained in long-term memory (LTM). The dominance of static features in the object token is what enables the same object to be identified in different contexts, whether it is moving or still.

Episodic memories group multiple object tokens together in contexts that capture the structure and setting of remembered events (reviewed by Eichenbaum, Yonelinas & Ranganath, 2007; Yonelinas *et al*., 2010; Zimmer & Ecker, 2010). Reactivation of an object token from LTM reactivates context information from the episodic memories with which the object token is associated, as well as other object tokens from those episodic memories. In many cases this reactivation is fast enough to seem instantaneous, lagging the feeling of familiarity by only about 100 ms. In other cases, however, object token reactivation takes seconds to minutes, and may require the acquisition of additional information that triggers the retrieval of a relevant episodic memory (Eichenbaum, Yonelinas & Ranganath, 2007).

Human beings interact with the same objects over and over again throughout significant periods of their lifespans; consider, for example, interactions with one's parents, spouse, children, house and car. The presence of object tokens referring to the same individual object in multiple episodic memories immediately raises the question of updating. Object tokens are constructed from object files by adding categorical information: they capture the *current* features and motion of an observed object. These current features are stored in episodic memories; otherwise one would not be able to recall incidental features such as what someone was wearing or how they behaved on a given occasion. Incidental features must be represented in a way that does not allow them to block re-identification of an object when it is encountered in a different context; it is the inability to block re-identification that renders such features incidental. Humans appear to resolve this representational issue by linking object tokens representing a given individual to a “singular category” containing only features that *do* block re-identification if not matched and are, therefore, effectively essential (Rips, Blok & Newman, 2006; Bullot & Rysiew, 2007; Xu, 2007; Bullot, 2009; Nichols & Bruno, 2010). Even the “essential” features contained in singular categories must, however, occasionally be updated; a mis-matched apparent age can block re-identification of a child, for example, but apparent age must be periodically updated or it will block re-identification permanently and hence incorrectly. Over extended time, therefore, singular categories must be represented as adjustable “models” of an individual that capture historically- as well as currently-essential features.

No stringency of feature-based recognition, division between incidental and essential features, or feature-updating scheme is, however, able fully to account for human object re-identification capabilities. When faced with significant feature change or indistinguishably-featured competitors, humans also employ causal constraints to determine whether a currently-perceived object is the temporal continuation of a previously-encountered object (reviewed by Rips, Blok & Newman, 2006; Scholl, 2007). Studies using change-blindness paradigms (reviewed by Simons & Ambinder, 2005) show that causal constraints on individual identity are strong enough to sometimes over-ride continuity of even ordinarily-essential features. Simons and Levin (1998), for example, conducted experiments in which approximately half of the subjects tested did not detect a change in conversation partner during the course of a brief interaction in broad daylight, even though the two conversation partners had different facial features and significantly different attire. Causal constraints also affect episodic memory recall, as demonstrated by false-memory studies (reviewed by Mitchell & Johnson, 2000;

Henkel & Carbuto, 2008). Mendelsohn *et al.* (2009), for example, demonstrate increasing replacement over time of autobiographical memories by causally-plausible confabulations in a healthy, fully-functional individual. The history that is “filled in” in cases of change blindness or memory confabulation has clearly not been observed; it is *fictive* history generated to plausibly connect the observed present to the past of episodic memories (Fields, 2012a).

It is commonplace to think of both property updating and reasoning from causal constraints as implemented by heuristics that are either part of general world knowledge (Fodor, 2000; 2008) or embedded in domain-specific modules (Pinker, 1997) or mental models (Gentner, 2002). The lineage of AI programs from GPS (Newell & Simon, 1963) to CYC (Lenat & Guha, 1990) exemplifies the former way of thinking, while the lineage from expert systems (Hayes-Roth, Waterman & Lenat, 1983) to model-based planners (Kolodner, 1992) exemplifies the second. Alternative AI system architectures offer both a choice of heuristics to implement and a choice of data structures and operations with which to implement them; as noted previously, they also provide the built-in, typically Leibniz's Law compliant semantics of the chosen programming language. In the case of human problem solvers, both the heuristics that are employed and the data structures and operations that implement them must be discovered by experimentation, and no semantics beyond that implicit in the verbal or behavioral reports of the experimental subjects is given in advance. Despite these limitations, functional imaging studies as well as investigations of the consequences of lesions and other insults have now provided sufficient data to construct reasonably detailed neurofunctional models of both categorization and the deployment of causal constraints in human object re-identification.

How humans visually re-identify objects: Implementation

The characterized components of the human visual object tracking, categorization and identification system are summarized in Fig. 1. Visual input is processed by two mutually-modulatory pathways, the “ventral stream” through visual area V4 that processes feature information such as shape and color, and the “dorsal stream” through the medial temporal area (MT) that processes trajectory information (reviewed by Rizzolatti & Matelli, 2003; Nassi & Callaway, 2009). The ventral stream provides visual input to categorized object representations in lateral (LFG) and medial (MFG) fusiform gyrus for objects categorized as agents and non-agents, respectively (reviewed by Martin, 2007; Mahon & Caramazza, 2009). Categorization information flows “downward” from the anterior temporal pole (ATP), which appears to function in the integration of multi-modal semantic information (reviewed by Visser, Jefferies & Lambon Ralph, 2009; Kiefer & Pulvermüller, 2012), via the perirhinal cortex (PRC), which implements the categorized object token (Zimmer & Ecker, 2010). This feature-oriented processing stream provides “what” information for the encoding of episodic memories by the hippocampus (HC; reviewed by Ranganath, 2010).

*****

Fig. 1 about here.

*****

The dorsal, trajectory-oriented processing stream provides input to representations of the motions of categorized agents and non-agents in superior temporal sulcus (STS) and medial temporal gyrus (MTG) respectively (Martin, 2007; Mahon & Caramazza, 2009), as well as input to posterior-parietal areas involved in the control of actions (Rizzolatti & Matelli, 2003; Nassi & Callaway, 2009; see also below). These motion-based representations provide "where" information to HC via the parahippocampal cortex (PHC), which organizes the spatial layout or context of ongoing events (Eichenbaum, Yonelinas & Ranganath, 2007; Ranganath, 2010). Experiments that probe the criteria for within-scene object persistence at the object file level appear to be accessing transitory trajectory representations across this MT-to-PHC network (Fields, 2011a).

The representations in both PHC and PRC are reinstated when episodic memories are recalled; these reinstated representations compete with and modulate current, perceptually-driven representations (Eichenbaum, Yonelinas & Ranganath, 2007; Zimmer & Ecker, 2010; Ranganath, 2010). As discussed above, however, the pathways shown in Fig. 1 do not encode sufficient information to re-identify currently-perceived objects as being the same individuals as recalled participants in remembered episodes. At best, they can recognize whether a current object is featurally-indistinguishable from a remembered object, and whether the current context is spatially-indistinguishable from a remembered context. If object features change significantly between contexts or a familiar object is encountered in an unfamiliar context, object re-identification can fail. Because the network shown in Fig. 1 encodes no model of either featural or contextual change, it does not implement a solution to the frame problem as it applies to object re-identification.

The human encoding of on-going events does not, however, involve only "what" and "where" information; it also involves tracking "how" and "why" objects came to be where they are in the context of active goals. This "event file" (Hommel, 2004) level of representation combines information from multiple sensory modalities, binds goals and actions to objects (Hommel, 2007; Spapé & Hommel, 2010), includes multiple levels of categorical abstraction (Colzato, Raffone & Hommel, 2006), and incorporates emotional responses to objects and events (Colzato, van Wouwe & Hommel, 2007). Goals, actions, objects, categories and emotions are represented in different areas of the brain; event files are, therefore, long-range patterns of correlated neuronal activity. Unimodal, e.g. visual event files are bound in 240 to 280 ms (Zmigrod & Hommel, 2010), the same time-frame required for object categorization. Scenes containing localized, categorized objects are accessible to consciousness after approximately 270 ms (Sergent, Baillet & Dehaene, 2005), suggesting that the event-file level of correlated neuronal activity corresponds to the "global workspace" that has been proposed as the basic substrate of conscious awareness and attentional control (Baars, 1997; Dehaene & Naccache, 2001; Dehaene & Changeaux, 2004). Episodic memory retrieval consistently reactivates not just the medial temporal lobe comprising HC, PRC and PHC but also the extended fronto-parietal attention control system (Wagner, 2005; Cabeza et al., 2008; Uncapher & Wagner, 2009; Ranganath, 2010), consistent with the incorporation of motivating goals and "how" and "why" information about object positions into experienced episodic memories, and hence with the reactivation of entire event files and the theoretical identification of the event-file level of representation with the global workspace.

By identifying the selection of relevant information as an architecturally-supported function of the global workspace, Shanahan and Baars (2005) provide a neurofunctionally-supported resolution of one aspect of the frame problem: the apparent isotropy of belief updating. Evidence increasingly supports the view that attention is feedback amplification of those patterns of network activity that incorporate current goal representations (reviewed by Chun, Golomb & Turk-Browne, 2011; Smallwood *et al.*, 2012). In any given problem-solving context, the representations that will be treated as "relevant" are those that most overlap with current goals. Which representations most overlap with current goals is determined largely by experience as represented in episodic memories; to the extent that these memories are missing information needed for solving the current problem, or that by absorbing the available attentional amplification they suppress retrieval of in-fact relevant information, problem solving will fail. The correlation between significant insights and relaxed attention noted by investigators from Archimedes onward, and confirmed by functional imaging studies (Kounios *et al*., 2006; 2007; Kounios & Beeman, 2009), reflects the tendency of attention to amplify information that has been considered relevant in the past at the expense of information that is actually relevant in the present. Allocating attention to memories previously associated with similar goals is, therefore, not an optimal mechanism for isotropy. It is, however, perhaps the only approximately rational mechanism available to organisms with small working memories and highly-constrained computational resources.

Resolution of the issue of isotropy by the mechanism of goal-driven attentional amplification does not, however, explain the human ability to re-identify objects as individuals. As noted above, this ability depends on the deployment of category-specific causal constraints that specify “how” and “why” objects move between contexts, and “how” and “why” their features can change as they do so. Humans incorporate two representations of “how” contextual and featural changes occur, both implemented by posterior-parietal components of the “mirror neuron system” that respond to observations of actions or motions carried out by others by activating pre-motor representations of such actions or motions as they would be carried out by the observer (reviewed by Rizzolatti & Craighero, 2004; Cattaneo & Rizzolatti, 2009; see Gazzola & Keysers, 2009 for high-resolution data). The responses of mirror-system components are re-configurable by experience (Catmur, Walsh & Heyes, 2007; Catmur *et al.*, 2008; Heyes, 2010), permitting non-agent motions to be re-represented as agent (i.e. self) motions (Schubotz & von Cramon, 2004; Engel *et al.*, 2007). As shown in Fig. 2, actions of identified autonomous agents are represented primarily within the inferior parietal lobule (IPL), including the agency-attribution system of the temporal-parietal junction (TPJ), with a bias in agent representation to the right hemisphere (Cattaneo & Rizzolatti, 2009). Movements of non-agents that are caused by agents and movements of self-propelled non-agents are represented primarily by mirror-system components of superior parietal lobule (SPL), with a bias toward the left hemisphere (Martin, 2007; Mahon & Caramazza, 2009). These representations are tightly coupled to goal-driven attentional reorienting systems (reviewed by Corbetta, Patel & Shulman, 2008).

*****
Fig. 2 about here.
*****

The primary human representation of “why” featural and contextual changes occur is that encoded by IPL: changes occur because agents act intentionally. This bias toward agency is well-documented both experimentally (reviewed by Scholl & Tremoulet, 2000) and culturally (reviewed by Bloom, 2007; Boyer & Bergstrom, 2008; Rosset, 2008). An event that results in strong activation of SPL but only weak activation of IPL is experienced as not caused by any agent; the lack of “why” information for such an event signals an event categorization conflict to the anterior cingulate. Conflict signals from anterior cingulate trigger goal revisions in prefrontal cortex (reviewed by Carter & van Veen, 2007; Rolls & Grabenhorst, 2008); revised goals drive attentional re-orientation, the retrieval of alternative episodic memories and the construction of alternative FCHs. Individuals who experience events as not caused by agents and are motivated by the resulting conflict to search for “hidden” causes are termed “systemizers” (Baron-Cohen, 2002; 2008) or “mechanizers” (Crespi & Badcock, 2008); natural scientists, technologists, engineers and mathematicians tend to be systemizers. The construction by systemizers of “why” information for events not caused by agents appears to involve the simultaneous activation of pre-motor representations of force-delivering actions (Fields, 2012b) and suppression of self-as-agent representations (Fields, 2011b). Children typically develop an ability to construct such representations around age four (Sobel *et al.,* 2007), well after they develop the ability to attribute intentions to agents (Saxe, Carey & Kanwisher, 2004), the tendency to interpret the behavior of inanimate objects as intentional (Boyer, Pan & Bertenthal, 2011; Cicchino, Aslin & Rakison, 2011), and the tendency to regard inanimate objects as products of intentional design and construction (Kelemen, 2004).

Identifying a perceived object as the causal continuer of a remembered object requires generating “how” information connecting the remembered object and context to the current object and context. Because it is generated by the mirror neuron system, this information is represented by a pre-motor plan for a bodily motion that is both observable and executable. Pre-motor plans specify applied forces that supply both intentional and mechanical “why” information; as noted above, the representation of purely mechanical “why” information requires the suppression of the intentional component of the pre-motor plan. If viewed formally, the inferences that generate these pre-motor plans are structure mappings (Gentner, 1983) that preserve force-motion relations (Fields, 2011c, 2012b); when executed simultaneously at multiple levels of abstraction, such pre-motor inferences appear to be capable of supporting general action planning (Schubotz, 2007; Bubic, von Cramon & Schubotz, 2010). The “how” and “why” information specified by such a pre-motor plan implicitly locates a re-identified object in space during the period of non-observation between the remembered context and the present; it also attributes particular causes to each of the object's movements and feature changes. Such plans are, therefore, complete – even if often quite abstract – fictive causal histories (FCHs) of the objects that they re-identify. As object re-identification in most instances occurs within the few hundred ms required for consciousness at the event-file level of representation – think of re-identifying your spouse or your car – the FCH linking a retrieved episodic memory to the current event file must be generated within this timeframe; the effects of prefrontal goal representations and even emotions on this process are, therefore, top-down without being consciously “deliberative”. Even in cases in which re-identification is significantly delayed, people tend to realize suddenly that a perceived object is the same thing as one previously encountered, as opposed to consciously inferring it from deliberatively-considered evidence (Eichenbaum, Yonelinas & Ranganath, 2007), indicating that FCH construction is fast and automatic even when it is delayed.

The FCHs that humans construct in the course of re-identifying objects are based only on the data that are available: the currently-perceived situation, the retrieved episodic memory, and the singular and general categorical constraints associated with its component object tokens. Such data clearly under-determine actual history; FCHs are, therefore, by their construction rough approximations at best of reality. Nothing guarantees even approximate accuracy. The resolution of FCH construction is, moreover, limited by the motion representations available to SPL. “Someone put it there” may be as much detail as can be generated, but may nonetheless be sufficient to re-identify – accurately or otherwise – your missing umbrella. The rapidity of FCH construction and their lack of causal detail are consistent with the commonplace experience of “not knowing how something got there” even though it has been re-identified as the causal continuation of a familiar object.

The role of parietal cortex in generating the “how” and “why” information needed for object re-identification that is sketched here suggests an answer to one of the most puzzling questions concerning object persistence: the question of how human beings perceive themselves as persistent through time, and hence perceive their episodic memories as *their* memories. Posterior parietal systems, particularly

TPJ, are key components of the “default system” that maintains the sense of self (reviewed by Buckner, Andrews-Hanna & Schacter, 2007). The experienced consequences of parietal lesions indicate that pre-motor processes are involved in the experience of agency and hence of of self (reviewed by de Jong, 2011). It is at least plausible, therefore, that the activity of constructing FCHs that re-identify perceived objects is experienced, at least in part, also as a re-identification of the self as a persistent experiencer.

The human solution: Default causal knowledge and resource-limited isotropy

Together with the limitations on attentional resources imposed by the global workspace architecture, the construction of FCHs linking re-identified individual objects to retrieved episodic memories provides a general heuristic solution to the frame problem. This solution divides the objects in any perceived context into three classes: (1) those objects that are categorized by type but not re-identified as known individuals; (2) those objects that are re-identified as known individuals, and for which neither features nor location within the context have changed since the most recent episodic memory; and (3) those objects that are re-identified as known individuals, but for which features, location within the context, or both have changed since the most recent episodic memory. The frame problem clearly does not arise for objects in the first class: if they are not causal continuations of objects present to observation with sufficient salience to be recorded as episodic-memory components at some time prior to taking an action, their individual properties are not available for updating in consequence of that action. The presence of such objects in a post-action context is irrelevant if they are not salient, and surprising if they are. The frame problem also does apply to objects in the second class; these are the “sleeping dogs” (McDermott, 1987) for which a default FCH of “nothing happened” is assumed to be accurate. It is for objects in the third class that the frame problem poses a significant issue: these are the objects for which the FCH constructed to re-identify them has significant causal content. For objects in this third class, the constructed FCH accounts for the featural and contextual changes from the remembered to the current context that are evident to observation: such an accounting is precisely what is required for object re-identification by FCH construction.

Predicting the consequences of a deliberately planned action requires determining which objects in the current, pre-action context will be “sleeping dogs” for which no updating is required. Doing this requires imagining the action's outcome using some set of sensory modalities, in most cases a combination of vision and spoken language. Available data indicate that imaginations are event files constructed top-down; up to differences in modulation by the prefrontal “reality monitoring” system (Simons *et al*., 2008), they activate the same post-detection modal-network components that are activated when event files are constructed bottom-up by perceptual processes (reviewed by King, 2006; Kosslyn, Thompson & Ganis, 2006; Moulton & Kosslyn, 2009; Hubbard, 2010). When a possible outcome of an action is imagined, the question of which objects are sleeping dogs is resolved by the process of constructing the corresponding event file: the sleeping dogs are the objects connected to the pre-action context by trivial FCHs. The objects that are not sleeping dogs are mapped into the imagined outcome by non-trivial FCHs that specify, by explicit construction, how the action or its predictable side-effects moved or changed them. The range of scenarios that can be imagined as outcomes of a given object-directed action in a given context is, therefore, limited by the same categorical constraints and action/motion representations that limit FCH construction; hence the range of outcomes that can be imagined is no larger than the range of actual situations in which all of the participating objects would be re-identifiable as persistent individuals. Solutions to the object re-identification problem thus bound solutions to the frame problem as a whole. In the case of human problem solvers, therefore, the problem of object re-identification is not merely an instance of the frame problem: for humans, the problem of object re-identification *is* the frame problem.

Viewing the frame problem as equivalent, for humans, to the problem of object re-identification allows the limitations expected in the human solution of the frame problem, which have in general not been investigated experimentally, to be read off from limitations in the human solution to the object re-identification problem that have been investigated in considerable detail. As discussed above, humans are highly biased in their allocation of attention and are susceptible to change blindness under easily-manipulable conditions. Human episodic memories are strongly coupled to emotional responses and are susceptible to emotional manipulation as well as to cross-talk between similar contexts containing different objects. One can expect, therefore, that human beings will make systematic updating errors following actions in cases involving overly-narrow attentional focus, overly-optimistic assumptions of situational familiarity, emotional stress, and contextual ambiguity. All of these are situations in which the problems caused by unintended side-effects, i.e. incorrect frame-problem solutions, are all too familiar. The quality of frame-problem solutions can, moreover, be expected to be highly dependent on expertise. Accurate categorical constraints on the behavior of objects – inanimate or animate – together with highly-developed pre-motor simulation capabilities can be expected to yield better frame problem solutions; deficits in either of these areas can be expected to yield errors. Reliance on purely agent-driven causal models, in particular, will produce inaccurate frame problem solutions whenever the intentions assigned to the "agent" either under- or over-constrain the behavior of objects relative to the actual causal processes operating in the situation.

Conclusion

The frame problem has traditionally been considered independently of the object re-identification problem, which with typical assumptions about programming-language semantics does not arise. What has been shown here is that, if the model of object re-identification by FCH construction is correct, the frame problem and the problem of object re-identification are equivalent for the human cognitive architecture. If this model is correct, the human solution of the frame problem is clearly "embodied" (Gallese & Lakoff, 2005; Barsalou, 2008; Kiefer & Pulvermüller, 2012): it involves causal inferences implemented by the pre-motor system. It is, however, equally clearly algorithmic: the data available for processing are highly constrained, and the procedures that process them are architecturally fixed.

Viewed through the abstracting lens of symbolic data structures, the human solution to the frame problem is a spaghetti-code kluge of barely-informative heuristics, case-based reasoning and theoretical modeling, all overlaid with an experience-dependent hodge-podge of emotional and attentional biases. What saves this system from functional irrelevance is that human episodic memory retrieval is controlled not only by the similarity machine implemented by temporal-lobe categorization, but also by the analogy machine of the parietal-lobe pre-motor action-planning system, which computes structure mappings between alternative causal histories (Schubotz, 2007; Bubic, von Cramon & Schubotz, 2010; Fields, 2011c; 2012b). The frame problem was originally motivated by a view of problem solving as fundamentally deductive, but it has long been clear that workable solutions are abductive. In the case of human beings, we can now claim a moderately good understanding of how this abductive process works.

Acknowledgment

Thanks to Eric Dietrich for almost 30 years of discussions around the frame problem.

Conflict of interest statement

The author states that he has no financial or other conflicts of interest relevant to the research reported.

References

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

Baars, B. J. (1997). In the theater of consciousness: Global workspace theory, a rigorous scientific theory of consciousness. *Journal of Consciousness Studies 4*, 292-309.

Baron-Cohen, S. (2002). The extreme male brain theory of autism. *Trends in Cognitive Sciences 2,* 248–254.

Baron-Cohen, S. (2008). Autism, hypersystemizing, and truth. *Quartarly Journal of Experimental Psychology, 61*(1), 64–75.

Barsalou, L. W. (2008). Grounded cognition. *Annual Review of Psychology 59*, 617-645

Bloom, P. (2007). Religion is natural. *Developmental Science 10*(1), 147-151.

Boyer, P. & Bergstrom, B. (2008). Evolutionary perspectives on religion. *Annual Review of Anthropology* 37, 111–130.

Boyer, T. W., Pan, J. S. & Bertenthal, B. I. (2011). Infants' understanding of actions performed by mechanical devices. *Cognition 121*, 1-11.

Bubic, A., von Cramon, D. Y. & Schubotz, R. I. (2010). Prediction, cognition and the brain. *Frontiers in Psychology – Human Neuroscience 4*, 25 (DOI: 10.3389/fnhum.2010.00025).

Buckner, R. L., Andrews-Hanna, J. R. & Schacter, D. L. (2007). The brain's default network: Anatomy, function, and relevance to disease. *Annals of the New York Academy of Sciences* 1124, 1-38.

Bullot, N. (2009). Toward a theory of the empirical tracking of individuals: Cognitive flexibility and the functions of attention in integrated tracking. *Philosophical Psychology 22*(3), 353-387.

Bullot, N. J. & Rysiew, P. (2007). A study in the cognition of individuals' identity: Solving the problem of singular cognition in object and agent tracking. *Consciousness and Cognition 16*, 276-293.

Burke, L. (1952). On the tunnel effect. *Quartarly Journal of Experimental Psychology 4*, 121–138.

Cabeza, R., Ciaramelli, E., Olson, I. R. & Moscovitch, M. (2008). Parietal cortex and episodic memory: An attentional account. *Nature Reviews Neuroscience 9*(8), 613-625.

Carter, C. S., & van Veen, V. (2007). Anterior cingulate cortex and conflict detection: An update of theory and data. *Cognitive, Affective & Behavioral Neuroscience 7*(4), 367–379.

Catmur, C., Gillmeister, H., Bird, G., Liepelt, R., Brass, M., and Heyes, C. (2008). Through the looking glass: Counter-mirror activation following incompatible sensorimotor learning. *European Journal of Neuroscience 28*, 1208–1215.

Catmur, C., Walsh, V., and Heyes, C. (2007). Sensorimotor learning configures the human mirror system. *Current Biology 17*, 1527–1531.

Cattaneo, L., and Rizzolatti, G. (2009). The mirror neuron system. *Archives of Neurology 66*, 557–560.

Chun, M. M., Golomb, J. D. & Turk-Browne, N. B. (2011). A taxonomy of external and internal attention. *Annual Review of Psychology 62*, 73-101.

Cicchino, J. B., Aslin, R. N. & Rakison, D. H. (2011). Correspondences between what infants see and know about causal and self-propelled motion. *Cognition 118*(2), 171-192.

Colzato, L., Raffone, A. & Hommel, B. (2006). What do we learn from binding features? Evidence for multilevel feature integration. *Journal of Experimental Psychology: Human Perception and Performance 32*, 705-716.

Colzato, L. S., van Wouwe, N. C. & Hommel, B. (2007). Feature binding and affect: Emotional modulation of visuo-motor integration. *Neuropsychologia 45*, 440-446.

Corbetta, M., Patel, G. & Shulman, G. L. (2008). The reorienting system of the human brain: From environment to theory of mind. *Neuron 58*(3), 306-324.

Crespi, B., & Badcock, C. (2008). Psychosis and autism as diametrical disorders of the social brain. *Behavioral and Brain Sciences 31*, 241–320.

Decety, J. & Lamm, C. (2007). The role of the right temporoparietal junction in social interaction: How low-level computational processes contribute to meta-cognition. *The Neuroscientist 13*(6), 580-593.

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

Dehaene, S. & Changeaux, J.-P. (2004). Neural mechanisms for access to consciousness. In Gazzaniga, M. (Ed.) *The Cognitive Neurosciences*, 3rd Ed (pp. 1145-1157). New York, Norton.

de Jong, B. M. (2011). Neurology of widely embedded free will. *Cortex* 47, 1160-1165.

Eichenbaum, H., Yonelinas, A. R. & Ranganath, C. (2007). The medial temporal lobe and recognition memory. *Annual Review of Neuroscience 30*, 123-152.

Engel, A., Burke, M., Fiehler, K., Bien, S., & Rosler, F. (2007). How moving objects become animated: The human mirror system assimilates non-biological movement patterns. *Social Neuroscience 3*, 368–387

Fields, C. (2011a). Trajectory recognition as the basis for object individuation: A functional model of object file instantiation and object-token encoding. *Frontiers in Psychology – Perception Science 2*, 49 (DOI: 10.3389/fpsyg.2011.00049).

Fields, C. (2011b). From ‘‘Oh, OK’’ to ‘‘Ah, yes’’ to ‘‘Aha!’’: Hyper-systemizing and the rewards of insight. *Personality and Individual Differences 50*, 1159-1167.

Fields, C. (2011c). Implementation of structure-mapping inference by event-file binding and action planning: A model of tool-improvisation analogies. *Psychological Research 75*, 129–142

Fields, C. (2012a). The very same thing: Extending the object token concept to incorporate causal constraints on individual identity. *Advances in Cognitive Psychology* 8, 234-247.

Fields, C. (2012b). Motion as manipulation: Implementation of force–motion analogies by event-file binding and action planning. *Cognitive Processing* 13, 231-241.

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

Fodor, J. A. (2000). *The Mind Doesn't Work That Way: The Scope and Limits of Computational Psychology.* MIT Press.

Fodor, J. A. (2008). *The Language of Thought Revisited*. Oxford University Press.

Ford, K. M. & Pylyshyn, Z. W. (Eds) (1996). *The Robot's Dilemma Revisited: The Frame Problem in Artificial Intelligence*. Norwood, NJ: Ablex.

Gallese, V. & Lakoff, G. (2005). The brain’s concepts: The role of sensory- motor systems in conceptual knowledge. *Cognitive Neuropsychology 22*, 455–479

Gao, T., and Scholl, B. J. (2010). Are objects required for object files? Roles of segmentation and spatiotemporal continuity in computing object persistence. *Visual Cognition 18*, 82–109.

Gazzola, V. & Keysers, C. (2009). The observation and execution of actions share motor and somatosensory voxels in all tested subjects: Single-subject analyses of unsmoothed fMRI data. *Cerebral Cortex 19*, 1239–1255.

Gentner, D. (1983). Structure-mapping: a theoretical framework for analogy. *Cognitive Science 7*, 155–170.

Gentner, D. (2002). Mental models, Psychology of. In: N. J. Smelser & P. B. Bates (Eds.) *International Encyclopedia of the Social and Behavioral Sciences*. Amsterdam: Elsevier (pp. 9683-9687).

Gerhardstein, P., Schroff, G., Dickerson, K., & Adler, S. A. (2009). The development of object recognition through infancy. In B. C. Glenyn & R. P. Zini (Eds) *New Directions in Developmental Psychobiology* (pp. 79–115).

Hayes-Roth, F., Waterman, D. & Lenat, D. (1983). An overview of expert systems. In F. Hayes-Roth, D. Waterman & D. Lenat (Eds) *Building Expert Systems*. Reading, MA: Addison Wesley (pp. 3-29).

Henkel, L. A. & Carbuto, M. (2008). Remembering what we did: How source misattributions arise from verbalization, mental imagery, and pictures. In: Kelley, M. R. (Ed.) *Applied Memory* (pp. 213-234). Nova Science Publishers.

Heyes, C. (2010). Where do mirror neurons come from? *Neuroscience and Biobehavioral Reviews 34*(4), 575-583.

Hollingworth, A., and Franconeri, S. L. (2009). Object correspondence across brief occlusion is established on the basis of both spatiotemporal and surface feature cues. *Cognition 113*, 150–166.

Hommel, B. (2004). Event files: Feature binding in and across perception and action. *Trends in Cognitive Sciences 8 (11),* 494-500.

Hommel, B. (2007). Feature integration across perception and action: Event files affect response choice. *Psychological Research 71*, 42-63.

Hubbard, T. L. (2010). Auditory imagery: Empirical findings. *Psychological Bulletin 136*(2), 302-329.

Kelemen, D. (2004). Are children “intuitive theists”? Reasoning about purpose and design in nature. *Psychological Science 15*(5), 295-301.

Kiefer, M. & Pulvermüller, F. (2012). Conceptual representations in mind and brain: Theoretical developments, current evidence and future directions. *Cortex* 48, 805-825.

King, A. J. (2006). Auditory neuroscience: Activating the cortex without sound. *Current Biology 16*(11), R410-R411.

Kolodner, J. L. (1992). An introduction to case-based reasoning. *Artificial Intelligence Review 6*, 3-34.

Kosslyn, S. M., Thompson, W.L. & Ganis, G. (2006). *The Case For Mental Imagery*. Oxford University Press, New York

Kounios, J., & Beeman, M. (2009). The ‘‘Aha!’’ moment: The cognitive neuroscience of insight. *Current Directions in Psychological Science 18*(4), 210–216.

Kounios, J., Fleck, J. I., Green, D. L., Payne, L., Stevenson, J. L., Bowden, E. M., *et al.* (2007). The origins of insight in resting-state brain activity. *Neuropsychologia, 46,* 281–291.

Kounios, J., Frymiare, J. L., Bowden, E. M., Fleck, J. I., Subramaniam, K., Parrish, T. B., *et al*. (2006). The prepared mind: Neural activity prior to problem presentation predicts subsequent solution by sudden insight. *Psychological Science 17,* 882–890.

Lenat, D. & Guha, R. V. (1990). *Building Large Knowledge-Based Systems: Representation and Inference in the CYC Project*. Addison-Wesley.

Mahon, B. Z. & Caramazza, A. (2009). Concepts and categories: A cognitive neuropsychological perspective. *Annual Review of Psychology 60*, 27-51.

Martin, A. (2007). The representation of object concepts in the brain. *Annual Review of Psychology 58,* 25-45.

McCarthy, J. & Hayes, P. J. (1969), "Some philosophical problems from the standpoint of artificial intelligence". In: D. Michie & B. Meltzer (Eds.) *Machine Intelligence 4.* Edinburgh: Edinburgh University Press, pp. 463-502.

McDermott, D. (1987). We've been framed: Or, why AI is innocent of the frame problem. In: Z.W. Pylyshyn (Ed.), *The Robot's Dilemma: The Frame Problem in Artificial Intelligence*, Ablex Publishing, Norwood, NJ, (pp. 113–122).

Mendelsohn, A., Furman, O., Navon, I. & Dudai, Y. (2009). Subjective versus documented reality: A case study of long-term real-life autobiographical memory. *Learning & Memory 16,* 142-146.

Mitchell, K.J., & Johnson, M.K. (2000). Source monitoring: Attributing mental experiences. In E. Tulving & F.I.M. Craik (Eds.), *The Oxford Handbook of Memory*. New York: Oxford University Press (pp. 179-195).

Moulton, S. T. & Kosslyn, S. M. (2009). Imagining predictions: mental imagery as mental emulation. *Philosophical Transactions of the Royal Society B 364*, 1273–1280

Nassi, J. J., and Callaway, E. M. (2009). Parallel processing strategies of the primate visual system. *Nature Reviews Neuroscience 10,* 360–372.

Newell, A. & Simon, H. A. (1963). GPS: A program that simulates human thought. In E. A. Feigenbaum & J. Feldman (Eds.) *Computers and Thought.* New York: McGraw Hill (pp. 109-133).

Nichols, S. & Bruno, M. (2010). Intuitions about personal identity: An empirical study. *Philosophical Psychology 23*(3), 293-312.

Pinker, S. (1997). *How the Mind Works.* Norton.

Pylyshyn, Z. W. (Ed.) (1987). *The Robot's Dilemma: The Frame Problem in Artificial Intelligence*.

Norwood, NJ: Ablex.

Pylyshyn, Z. (2009). Perception, representation, and the world: The FINST that binds. In: D. Dedrick & L. Trick (Eds.) *Computation, Cognition, and Pylyshyn.* Cambridge, MA: MIT Press (pp. 3–48).

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.

Rips, L., Blok, S. & Newman, G. (2006). Tracing the identity of objects. *Psychological Review 133*(1), 1-30.

Rizzolatti, G., and Matelli, M. (2003). Two different streams form the dorsal visual system: Anatomy and functions. *Experimental Brain Research 153*, 146–157.

Rizzolatti, G., and Craighero, L. (2004). The mirror neuron system. *Annual Review of Neuroscience 27*, 169–192.

Rolls, E. T., & Grabenhorst, F. (2008). The orbitofrontal cortex and beyond: From affect to decision-making. *Progress in Neurobiology 86*, 216–244.

Rosset, E. (2008). It's no accident: Our bias for intentional explanations. *Cognition 108*, 771–780.

Samuels, R. (2010). Classical computationalism and the many problems of cognitive relevance. *Studies in the History and Philosophy of Science Part A 41*(3), 280-293.

Saxe, R., Carey, S., & Kanwisher, N. (2004). Understanding other minds: Linking developmental psychology and functional neuroimaging. *Annual Review of Psychology 55*, 87–124.

Scholl, B. J. (2007). Object persistence in philosophy and psychology. *Mind & Language 22*(5), 563-591.

Scholl, B., and Tremoulet, P. (2000). Perceptual causality and animacy. *Trends in Cognitive Sciences 4*, 299–309.

Schubotz, R. I. (2007). Prediction of external events with our motor system: Towards a new framework. *Trends in Cognitive Sciences 11*(5):, 211–218

Schubotz, R. I. & van Cramon, D. Y. (2004). Sequences of abstract nonbiological stimuli share ventral premotor cortex with action observations and imagery. *Journal of Neuroscience 24*(24), 5467–5474

Sergent, C., Baillet, S. & Dehaene, S. (2005). Timing of the brain events underlying access to consciousness during the attentional blink. *Nature Neuroscience 8*(10) 1391-1400.

Shanahan, M. P. & Baars, B. J. (2005), Applying global workspace theory to the frame problem. *Cognition 98*(2), 157–176.

Simons, D. J. & Ambinder, M. S. (2005). Change blindness: Theory and consequences. *Current Directions in Psychological Science 14*(1), 44-48.

Simons, D. J. & Levin, D. T. (1998). Failure to detect changes to people in a real-world interaction. *Psychonomic Bulletin & Review 5,* 644-649.

Simons, J. S., Henson, R. N. A., Gilbert, S. J. & Fletcher, P. C. (2008). Separable forms of reality monitoring supported by anterior prefrontal cortex. *Journal of Cognitive Neuroscience 20*(3), 447-457.

Smallwood, J., Brown, K., Baird, B. & Schooler, J. W. (2012). Cooperation between the default mode network and the frontal–parietal network in the production of an internal train of thought. *Brain Research 1428,* 60-70.

Sobel, D., Yoachim, C., Gopnik, A., Meltzoff, A., & Blumenthal, E. (2007). The blicket within: Preschooler’s inferences about insides and causes. *Journal of Cognitive Development*, *8*(2), 159–182.

Spapé, M. M. & Hommel, B. (2010). Actions travel with their objects: Evidence for dynamic event files. *Psychological Research 74,* 50-58.

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

Uncapher, M. & Wagner, A. D. (2009). Posterior parietal cortex and episodic encoding: Insights from fMRI, subsequent memory effects and dual-attention theory. *Neurobiology of Learning and Memory 91*(2), 139 (DOI: 10.1016/j.nlm.2008.10.011).

Visser, M., Jefferies, E. & Lambon Ralph, M. A. (2009). Processing in the anterior temporal lobes: A meta-analysis of the functional neuroimaging literature. *Journal of Cognitive Neuroscience 22*(6), 1083-1094.

Wagner, A. D., Shannon, B. J., Kahn, I. & Buckner, R. L. (2005). Parietal lobe contributions to episodic memory retrieval. *Trends in Cognitive Sciences 9*(9), 445-453.

Wheeler, M. (2008). Cognition in context: Phenomenology, situated cognition and the frame problem. *International Journal of Philosophical Studies 16*(3), 323-349.

Xu, F. (2007). Sortal concepts, object individuation, and language. *Trends in Cognitive Sciences 11*(9), 400-406.

Yonelinas, A. P., Aly, M., Wang, W.-C. & Koen, J. D. (2010). Recollection and familiarity: Examining controversial assumptions and new directions. *Hippocampus 20,* 1178-1194.

Zimmer, H. D. & Ecker, U. K. D. (2010). Remembering perceptual features unequally bound in object and episodic tokens: Neural mechanisms and their electrophysiological correlates. *Neuroscience and Biobehavioral Reviews 34*(7), 1066-1079.

Zmigrod, S. & Hommel, B. (2010). Temporal dynamics of unimodal and multimodal feature binding. *Attention, Perception and Psychophysics 72*(1), 142-152.

Figure captions

Fig. 1: Human implementation of object tracking, object categorization, and episodic memory of “what” and “where” information. Solid arrows show “bottom-up” perceptual processing and dashed arrows shown “top-down” categorization; the double arrow between ATP and HC indicates off-line category refinement. Lateral connections are suppressed for simplicity of presentation. Abbreviations for anatomic areas are as defined in the text.

Fig. 2: Human implementation of agency attribution, actions by agents, and motions of self-propelled non-agents. Solid arrows show pathways leading to episodic memory encoding by HC; double arrows indicate off-line category refinement (ATP – HC) and long-range connections within the fronto-parietal attention-control network. Feedback and lateral connections as well as hemispheric differences are suppressed for simplicity of presentation. Abbreviations for anatomic areas are as defined in the text.

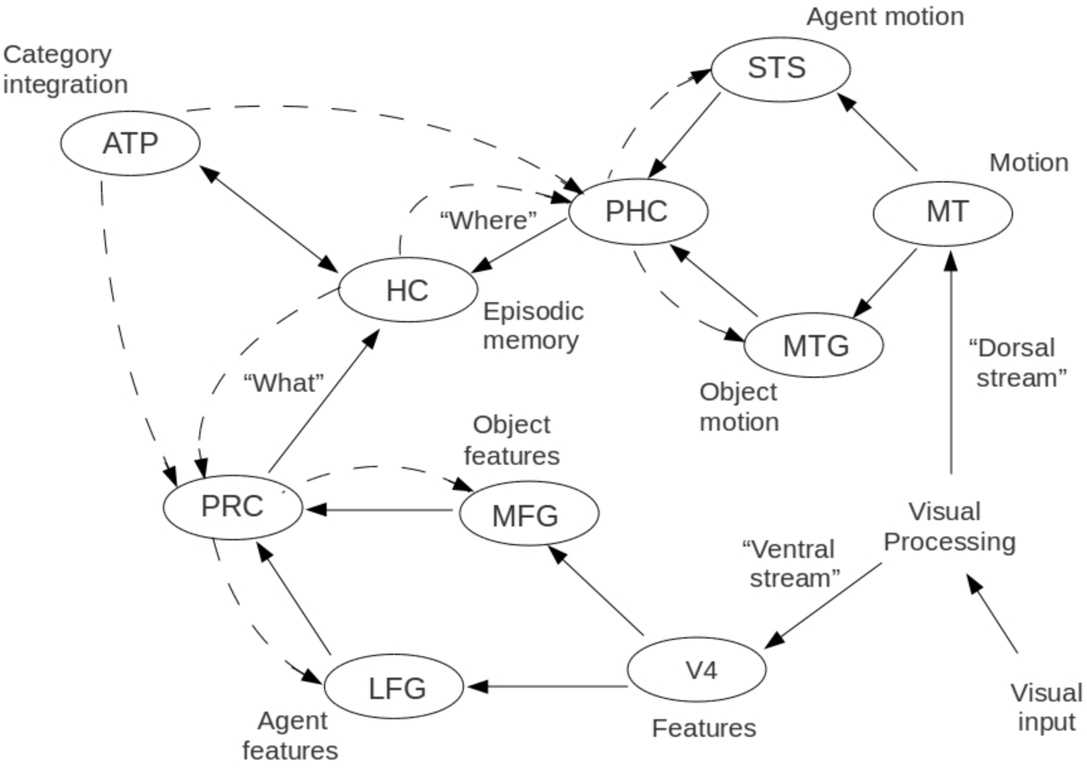
Agent motion
STS
Category integration
ATP
Motion
PHC
MT
"Where"
HC
Episodic memory
MTG
"Dorsal stream"
"What"
Object motion
Object features
PRC
MFG
Visual Processing
"Ventral stream"
V4
LFG
Visual input
Features
Agent features

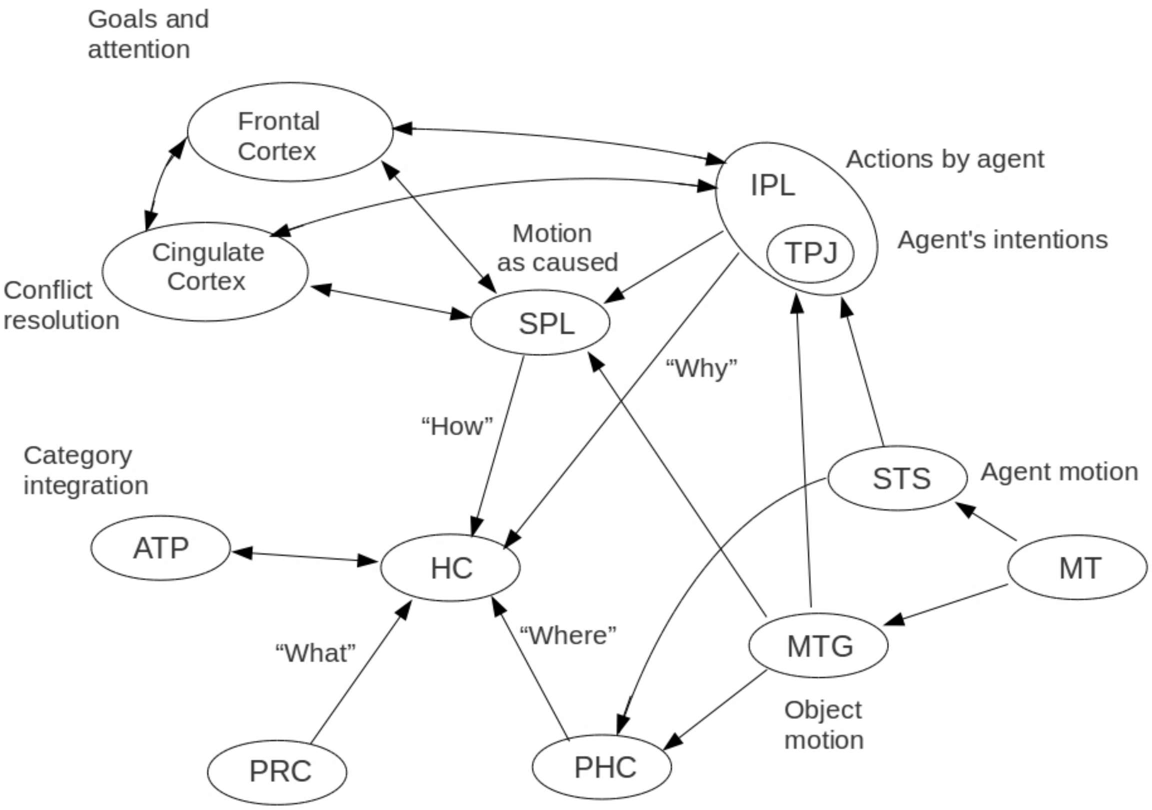
Goals and attention
Frontal Cortex
IPL
Actions by agent
TPJ
Agent's intentions
Cingulate Cortex
Conflict resolution
Motion as caused
SPL
"Why"
"How"
Category integration
STS
Agent motion
ATP
HC
MT
"What"
"Where"
MTG
Object motion
PRC
PHC