

Some Computational Desiderata for Recognizing and Reasoning About the Intentions of Others*

Paul F. Bello

Air Force Research Laboratory
Information Directorate
525 Brooks Rd.
Rome, New York 13323

Nicholas L. Cassimatis & Kyle McDonald

Department of Cognitive Science
Rensselaer Polytechnic Institute
Troy, New York 12180

Abstract

Reasoning about intentional action is a pervasive and critical skill in the human cognitive repertoire. Intentions have taken center-stage in discussions of how humans parse perceptual input, understand language, make moral judgments, and predict the behavior of conspecifics. In the quest to engineer machine intelligence, intentions have largely either been ignored entirely, or have been given oversimplified construals as either preference-orders over actions, or as simple predicates in computational theories of action. In this paper, we will motivate the need for intelligent systems capable of reasoning about the intentions of others by presenting a number of germane application areas, including those which deal with the integration of intention-recognition with other cognitive processes including dialogue processing. We then briefly review the relevant psychological literature on the development and operation of the human capacity to recognize and reason about the intentions of others. In doing so, we will extract a number of desiderata for the development of computational models of intention-recognition. We will then show how these requirements motivate principled design choices for the construction of intelligent systems. Finally, we will close with a brief description of Polyscheme, a computational cognitive architecture which we feel sufficiently addresses the computational challenges entailed by our desiderata.

Introduction

The human capacity for recognizing and reasoning about intentional action is critical for enabling effective interaction with the environment, most especially environments populated with other intentional agents. Judgments of intentional acts help us to cut through much of the perceptual clutter which would otherwise paralyze and relegate us to a life filled with filtering out unimportant percepts. Advances in the development of intelligent systems capable of recognizing and reasoning about intentions have far-reaching implications for human-computer interaction (HCI) to multi-

agent interactions under uncertainty. Perhaps it is instructive to start off with an example scenario demonstrating the use of intentional attributions in noisy and confusing environments. To illustrate, we take a simple example of two employees interacting at a restaurant:

Waitress 1 turns to see a customer with his coffee cup in his hand, waving it around slightly. She also sees a ham sandwich on his plate. She turns to waitress 2 and exclaims: “Get the ham sandwich some more joe.” Waitress 2 looks around the room near waitress 1, identifies the appropriate customer, and immediately serves him some more coffee.

This narrative illustrates the complexity inherent in navigating social situations. If we were to build computational models of all three scenario participants, we would need assertions to the effect that the customer believes that by shaking his empty cup of coffee to get waitress 1’s attention amounts to a signal that his cup is empty. We also must represent the fact that the customer assumes that the waitress believes his cup to still be full by default. We then must represent that waitress 1 believes that waitress 2 is currently not disposed and has the same default information about the status of the customer’s coffee cup until she indicates otherwise. Waitress 1 must believe that her speech acts to disambiguate among numerous customers in the diner and possible items to be delivered by referring to “the ham sandwich” and “joe” will be efficacious. Finally, waitress 2 must use the given information under uncertainty amongst the number of possible referents in the room and potential items to be delivered in order to make a more informed decision about how to act. While these assertions probably don’t exhaustively characterize the situation as described, they provide an illustrative example of the type of expressive inference coupled with search and belief-revision that is required to navigate such domains. It should be clear that this exchange is peppered with reference-disambiguation via discernment of intentions. Joe initiates an intentional act by waving his cup in the air to attract the server’s attention. Seemingly, Joe believes that waving his coffee cup rather than his fork indicates further that he’d like more coffee. In his selection of intentional acts to perform, he must adopt the (mental) perspective of the server, and best select an act which reduces the proverbial “signal-to-noise ratio,” conditioned on the as-

*The authors wish to thank Arthi Murugesan, Magdalena Bugajska, Scott Dugas and David Pizarro for their stimulating ideas on these issues.

Copyright © 2007, Association for the Advancement of Artificial Intelligence (www.aaai.org). All rights reserved.

sumed beliefs of the server whom he is signalling. In this case, waving around his plate may indicate that's he'd like his plate cleared from the table, or that he is finished with his meal and in a rush to get a check. In any case, Joe selects his cup as the appropriate signalling device. Waitress 1 acknowledges Joe's cup-waving, and initiates an intentional act directed at waitress 2, by uttering "get the ham sandwich some more joe." In doing so, she implicitly prunes the space of possible customers for waitress 2 to attend to by providing a more efficient taxonomic scheme for describing customers. Waitress 1's body position and line-of-sight also play a role in helping waitress 2 figure out the proper referent. Finally, waitress 2 can use knowledge of the specific tables waitress 1 is responsible for serving as a constraining variable in choosing the right action sequence. In any case, we see from this simple examples of social interaction that both the production and recognition of intentional acts helps to make sense of an otherwise confusing world. Intentions, while expressible by a proposition or a predicate, are determined by a host of other factors including physical cues, assumed background knowledge, and re-representation of problem spaces.

A Brief Tour of Some Empirical Findings

One of the most fruitful ways to study those factors which comprise judgments of intentionality by humans is under the microscope of cognitive development. Data collected on developing children has provided some boundaries by which we can circumscribe our theoretical accounts of how humans understand and react to intentional actions performed by others. Our understanding of the human intention-recognition engine has been further illuminated by several striking findings from social cognitive neuroscience, which we shall touch upon momentarily. Before we begin, it's useful to draw a conceptual distinction between intentions and so-called *intention-in-action* (Searle 1983). Detecting intention-in-action roughly corresponds to attributing intentionality to an observed motor sequence executed either by the self or another by appealing to common features which define agent-initiated action. Some common features of intention-in-action include the effector tracing a continuous spatiotemporal path towards the object to be manipulated, a judgment about the biological plausibility of such a path¹, and a number of other judgments which help the observing agent categorize the actor as human – especially skin texture, direction of eye-gaze, and orientation of the actor's body with respect to the object to be acted upon (Woodward, Sommerville, & Guajardo 2001). Intention-in-action must be conceptually cloven from the broader notion of recognizing intentions, since prior intentions associated with performing an action are generally unobservable and lack the kinds of physical cues associated with recognizing intention-in-action. It is worth noting that experimentation has been done in order to gain insight as to how humans parse per-

¹This judgment can be thought of in terms of detecting the *physical ability* of the actor to take the particular action under observation. More general judgments about ability also seem pervasive in both determining which actions to take, and in making higher-order attributions of culpability to other agents.

ceptual input when observing agents taking action in their environments. It was demonstrated in (Newtson & Engquist 1976) that independent human subjects tended to have a high degree of inter-rater agreement as to the beginnings and endings of intentional actions within videos. This result was further strengthened in (Baldwin & Baird 1999) by using infant looking times² as a clue to boundaries on how the perceived behavior stream is chunked up by the mind. For infants, there is a very strong correlation between habituation/dishabituation shifts and the boundaries on intentional actions defined by normal adult subjects who are made to watch the same videos as their infant counterparts. This suggests that the human cognitive apparatus is built for parsing action into goal-directed chunks from the earliest days of our lives.

Under the assumption that actions are represented as hierarchically decomposed goal-directed chunks within the human cognitive architecture, we now turn to describing how ascriptions of intentionality may be made when humans reason about other humans who are taking actions within the purview of their perceptual apparatus. Traditionally, perception and action have been studied in relative isolation, and were considered to be coded in the brain using incompatible neural codings. This view has been significantly undermined in recent years through the discovery of several so-called *mirror neuron* systems in Macaque monkeys which produces near-identical activations when the subject performs a particular action and observes another actor perform the same action (Rizzolatti *et al.* 1996). This suggests that there is a plausible common language shared by perception and action allowing for interchangeable judgments about self and other. The representational correspondence existing between self and other points to observation of our own motor sequences and their necessity in performing particular self-initiated intentional actions as clues to detecting similar intentions in the actions of others. Such an explanation is consistent with the set of constraints we have defined for recognizing intention-in-action: those being relatively smooth spatiotemporal path planning, biological plausibility of motion, and some aspects of detecting human-specific features in the actor. In short, as we learn to take intentional actions ourselves and as we begin to explore the biological features and limitations of our own effectors, we can simultaneously use these constraints in recognizing intentional action initiated by others. Further neuroscientific exploration has uncovered more detailed distinctions in the mirror neuron system which activate when we *imagine* the performance of intentional action. This poses the rather vexing question of why this happens in the absence of proprioceptive feedback or any sort of connection to our own motor programs. In these special cases,

²Often referred to as "violation of expectancy" (VoE) paradigms. These experimental designs allow researchers to effectively work with pre-verbal infants. In the case of action-parsing VoE provide us interesting clues as to how interruptions in intentional sequences of actions correspond to surprise on the part of an infant observer, and the boundaries which seem to define what we consider to be an "action" versus unintentional movements.

our problem arises in trying to distinguish between self-intentions and intentions in others. An interesting solution to this problem is suggested in (Jeannarod 2003) who suggests that there are non-overlapping representational spaces in the brain corresponding to self and other along with the overlapping section which helps us make judgments about intentionality during observed action. Further neuroscientific corroboration of this distinction has come in the form of identifying specific brain areas corresponding to initiation of first-person action versus observation of third person actions (Ruby & Decety 2001). This tidily resolves the dilemma introduced by situations in which we mentally represent a motor sequence (say as mental imagery), but do not actually physically execute it. These results strongly motivate an *agent-independent* representation of action in the brain, and *encapsulated mental spaces* within which inferences about actions can be performed before the action in question may actually be passed onto the motor system for execution. Agent-independent representations existing within encapsulated mental spaces comport well with large portions of the psychological literature on the representation of other mental states, such as beliefs as defined by the so-called *simulation theory* and its variants (Goldman 2006; Nichols & Stich 2003) by explaining how and why children (and adults) make default attributions of beliefs and associated plans to other agents with whom they've had no prior contact. It seems we routinely impute these mental states to others, assuming that other agents are "Like Me" (Meltzoff 2005). Such a general-purpose design allows us to circumvent a host of complexity issues introduced by trying to tag every belief-desire-intention-action schema with the name of an appropriate agent.

Our discussion of recognizing intention-in-action, while informative, doesn't provide us all of the conceptual machinery we require to produce models of the kinds of interactions we saw in our diner scenario. In these more complex situations, detection of intention-in-action is tightly weaved together with a broader notion of intentions driven by world-knowledge (i.e. one waitress' knowledge of what section of the diner her co-worker is responsible for), or attributions of beliefs and desires to other agents. In the case of understanding dialogue or responding to imperatives seem to rely on less well-defined descriptions of concepts like ability, or skill. A judgment to ascertain whether or not an agent was responsible for a particular outcome is largely contingent on whether or not we can classify the action which caused it as being intentional. Attributes corresponding to whether or not the performing agent had the requisite knowledge/beliefs, skill, and desire to perform the action will play more prominent roles in making these kinds of determinations.

The Desiderata

Now that we have taken a whirlwind tour of some important empirical findings on intention recognition and generation, we have some initial clues as to the kinds of computational structures we might use in modeling these processes.

Multi/Supramodal Representations

Based on the plethora of data presented in (Meltzoff 2005), and in (Jeannarod 2003), on both infant imitation and neuroimaging of first and third-person perspective-taking, it seems reasonable to conclude that perception and action are commonly coded for within a supramodal space as suggested by the Active Intramodal Matching hypothesis. From a purely engineering-oriented point of view, this is a convenient insight, suggesting that feature-based perceptual data used in image/speech processing representationally should co-exist with both proprioceptive feedback and rule-based knowledge (perhaps causal rules describing how perceptual access is required for situated epistemic inferences). Some sort of integrated hybrid intelligent architecture seems to be a natural candidate for developing intentional intelligent systems. However, tight integration between representations in such models is rare, and is a challenge needing to be overcome.

Encapsulated Mental Spaces

As suggested in (Jeannarod 2003; Goldman 2006; Nichols & Stich 2003), reasoning about others requires that a distinction must be made between self and others in order to avoid a number of rather counterintuitive entailments following from a self-other mapping within the human cognitive architecture. While mirror neurons provide firm neuroscientific foundations for perception/action congruence, mental simulations must be run in appropriately cordoned worlds so that imagined actions on the part of self or others are not accidentally sent to our motor system for execution. These issues are also salient in the discussion of the cognitive components of pretend-play and other kinds of counterfactual reasoning. All of this motivates agent-independent action schemata, which are representationally more compact than their agent-indexed counterparts.

Gaze/Point Understanding

Embodied intelligent systems in physical environments must be able to use physical cues provided by other agents in order to disambiguate references in speech and to determine boundaries on intentional actions (Newtson & Engquist 1976). As was illustrated in our "ham sandwich" example, understanding physical cues and their connections to the beliefs, desires, intentions and obligations of other agents in the environment seems to be critical for robust social interaction.

Epistemic Inference

Many problems in recognizing and reasoning about intentions involve the incorporation of prior knowledge about the target agent's beliefs. Many suggestions for how humans reason about the mental states of others spring from the suggestion that a mapping exists between self and other in the mind. Meltzoff and colleagues have presented striking evidence for this position, resulting in the "Like Me" hypothesis. To computationally implement the mapping in question, it would seem that some form of identity or similarity-based computation would need to be performed in order to map

relevant aspects of self to other and vice-versa. The computation would also need to be subserved by a filtering mechanism which allowed only relevant dimensions of self/other to be mapped. Finally a capacity to perform counterfactual reasoning is required to inhibit one's own perspective on the world when trying to reason about other agent's (possibly false) beliefs.

Open-world Probabilistic Inference

As we learn and develop, we develop causal rules linking up perception of our own motor processes to internal states such as desire. Nowhere is this more apparent than in the development of reaching and grasping behavior in infancy. However, when making inferences about reaching behaviors in other agents, we have no perceptual access to their desires, leaving open the possibility that an observed reaching behavior is caused by something other than a putative instrumental desire. In this case, we are left with trying to determine the likelihood that a particular reaching behavior is goal-directed, rather than the result of a random movement caused by a spasm or related to an ambiguous target object within physical proximity of the reach. Unfortunately, purely Bayesian representations of these kinds of computations are difficult (if not impossible) to perform, since every observation of a reach adds more random variables to the problem, making the inference to intentionality an open-world problem.

Hierarchical Representations

Learning about intentions within a dynamic physical environment populated with other agents often requires that we represent and reason about space, time and beliefs. Traditional logical and probabilistic approaches to inference in these kinds of situations (via SAT-solving/planning, Multi-Chain Monte-Carlo, etc) are extraordinarily computationally complex since canonical representations of space/time/belief are exponential in the number of objects either physically present or posited (in the case of beliefs). But apparently, humans represent actions using intermediary goal-directed descriptions (i.e. subgoals) which do not require that time or space be represented at the micro-scale. Also, simulation-theoretic accounts of reasoning about other minds along with relevant looking-time studies do not require that beliefs be representationally indexed to individual agents, eliminating the need for canonical epistemic representations.

Polyscheme

The mind is an integrated set of both cognitive and perceptual processes that underlies human behavior. Currently, these processes are typically characterized using qualitatively different computational frameworks including search, rule-based reasoning, Bayesian networks, category hierarchies, neural networks, and constraint graphs to name a few. Seamlessly integrating all of these mechanisms into a cohesive policy for collaboratively solving such a problem apparently requires the capacity to reason over multiple computational representations to build a mutually-reinforced solution. The Polyscheme cognitive architecture (Cassimatis

2002; 2005) is designed to integrate multiple computational mechanisms in solving these sorts of problems. Polyscheme consists of a number of specialists which maintain their own proprietary representations that communicate with one another during problem-solving through coordination via a cognitive focus of attention. The two guiding principles behind Polyscheme's integration scheme are the *common function principle* and the *multiple implementation principle*. The common function principle states that many reasoning and problem solving strategies can be composed of sequences of the same set of common functions. The multiple implementation principle states that each procedural unit can be executed using algorithms based on multiple representations.

The common function principle (CFP) arises from the observation that many of the modeling frameworks in computational cognitive science share the same underlying functions. A sampling of some of these basic functions include:

- Forward inference: given a set of propositions, infer propositions that follow from them.
- Subgoaling: given the goal of determining the status of proposition P, make subgoals of determining the truth values of propositions which either directly entail or falsify P.
- Representing alternate worlds: represent and make inferences about hypothetical states of the world.
- Identity matching: given a set of propositions about an object, find other objects which might be identical to it.
- Grounding: Given a set of variables, return an admissible set of objects which bind to the variables in question.

Using these basic principles, many of the most popular algorithms in computational cognitive science and AI can be constructed. For example, search can be described as a situation in which one is unsure about the truth value of proposition P. In order to arrive at the value, we represent the world in which P is true and the world in which P is false, performing forward inference in both worlds. If further uncertainty about P remains, repeat the process, adding the accreted results of the last inferences to the knowledge base associated with each world. Similarly, propagation in Bayesian networks often relies on a process of stochastic simulation that can be characterized as proportionally representing worlds corresponding to the likelihood of A as opposed to NOT A. By breaking down these algorithms into common functions, it becomes possible to embed something like stochastic simulation inside of the search process, or having belief revision performed on-line in the middle of inference about the location of objects in space.

The multiple implementation principle (MIP) suggests that each of these common functions can be executed by a host of qualitatively different computational representations including rules, feed-forward neural networks, and associative memories. For example, it's clear to see that rules can readily be employed for forward inference, can be operated upon to generate subgoals, and can be populated with predicates containing arguments which correspond to the pos-

sible world(s) in which the rule applies. Similarly, simple feedforward neural networks which have been trained can perform forward inference via simple propagation of inputs through the various layers of the network. Neural networks can generate subgoals by asking for the value of a particular input node, given a specified value for the output node(s). Characteristically, neural networks can be used to perform identity-matching via their natural application as classification engines. Since every Polyscheme specialist implements common functions through the MIP, it is possible for neural network, rule-based, Bayesian, constraint, and category-based specialists to be used in concert when performing complex, dynamic problem solving. Further, since we have shown that many of the predominant paradigms in computational cognitive science can be simulated by sequential applications of common functions implemented by multiple representation-indifferent specialists, we attain flexible problem solving involving multimodal percepts and intermediate representations. Integration of common function sequences is achieved by Polyscheme's cognitive focus of attention mechanism. Polyscheme's focus manager forces all specialists to focus on and report their opinions on the truth-value of a particular proposition P at every point during problem execution. The selection of this particular implementation of attention is motivated by the existence of processing interference in the Stroop effect (Stroop 1935), which suggests that multiple mental processes operate simultaneously (word and color recognition, for example). Visual attention has also been demonstrated as an integrative mechanism for inputs from multiple sensory modalities (Triesman & Gelade 1980). Polyscheme is based on the notion that just as the perceptual Stroop effect can be generalized to higher-level nonperceptual cognition, that integrative perceptual attention suggests the existence of a higher-level cognitive focus of attention that is the mind's principle integrative mechanism.

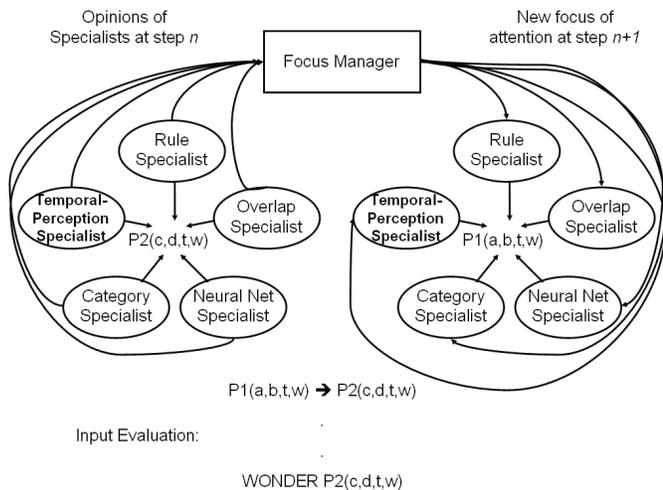


Figure 1: Focus Management in Subgoaling

As can be seen in Fig 1, during execution, if Polyscheme

desires to know the truth-value of a particular proposition such as P2, and has prior information to the effect that P2 is implied by P1, it will make a subgoal of finding out if P1 is true or not. Supposing this operation starts at execution step n, Polyscheme's focus manager requests that all specialists focus on P2, and report their opinions on its truth-value. But in this case, none of the specialists reach a consensus on P2. In fact, at least one specialist in this case (say the rule specialist) reports that in order for it to know the value of P2, it must first know if P1 is true. This is reported back to the focus manager as a request for attention on P1, and the process continues until consensus is reached among the specialists on P2, the original proposition in question. The general algorithm outlining this procedure roughly can be described in the following set of steps:

1. At each execution step n, Polyscheme's focus manager chooses a proposition to make the focus of attention.
2. Polyscheme collects the opinions of the specialists on the proposition being currently focused on.
3. Polyscheme reports these opinions to the other specialists.
4. Using their own computational mechanisms, the specialists process this new information, make inferences and decide which propositions they would help them make better information. These inferences define individual requests for attention by each specialist on the propositions they are individually most interested in.
5. The focus manager collects these requests for attention and decides the winner of the competition, which becomes the focus of attention at step n+1.

The scheme described above allows for *algorithmic opportunism*. Since we are able to reconstruct typical AI algorithms from our common functions, every computation performed by each specialist can be influenced and influence the inferences made by other specialists, enabling belief revision in the middle of search through spaces of alternative worlds, *et cetera*. While the specialists in Polyscheme were originally conceived of as the inferential building-blocks of infant physical reasoning (Cassimatis 2002), a number of reductions of higher-cognitive function down to the mechanisms of physical reasoning have been performed. Most notably, (Murugesan & Cassimatis 2006) demonstrated that a sophisticated syntax parser could be implemented on top of Polyscheme's physical reasoner, and (Bello & Cassimatis 2006b; 2006a) describes how cognitively plausible epistemic reasoning is a natural extension of mechanisms for physical inference. In light of our desiderata, Polyscheme seems to be a good first-step toward an intelligent system capable of recognizing and reasoning about intentions. Specifically, Polyscheme addresses the following:

- **Multi/Supramodal representations:** Polyscheme is specifically designed to reason over multiple representations within a common attention-driven framework. Specialists maintain their own proprietary representations, allowing for feature-based perceptual input to be processed by a neural network specialist which might need to be informed by a rule-based specialist.

- Encapsulated mental spaces: Every specialist in Polyscheme implements the “simulate alternative world” common function, which allows for hypothetical, future-hypothetical or counterfactual reasoning to be performed within the sandbox of each world. In the case of epistemic reasoning, encapsulated worlds represent situations in which Polyscheme “Is Like” the target agent to be reasoned about.
- Gaze/Point understanding: Polyscheme has been deployed on-board a robotic platform capable of reasoning about what others can and cannot see in the environment. Polyscheme is capable of making inferences about some of these unseen objects when referred to in dialogue by other agents. A full discussion can be found in (Cassimatis *et al.* 2004).
- Epistemic inference: Reasoning about other minds in Polyscheme is explained in detail in (Bello & Cassimatis 2006b). Polyscheme implements epistemic reasoning through the coordinated usage of alternate worlds, the identity specialist (to perform the mapping between self/other), and a form of inheritance of truth between propositions existing in different worlds allowing for reasoning about counterfactuals.
- Open-world probabilistic inference: Objects in Polyscheme can be lazily posited, and the system remains neutral on the truth-value of these objects until specifically asked to reason about them. Between lazy positing and lazy inference, Polyscheme has no need to represent the world canonically, drastically reducing the space-complexity in search.
- Hierarchical representation: Each proposition in Polyscheme is indexed with a time and a world in which it has a truth value. Polyscheme’s temporal reasoner is an implementation of Allen’s Temporal Interval Calculus, which allows for hierarchical representation of time. Rather than fluents needing to be represented canonically with respect to the time-points at which they may be true or false, we are able to use interval representations of time (i.e. Before, After) to eliminate large portions of the space of models to be evaluated.

Summary

To conclude, recognizing and reasoning about intentions in others entails a series of difficult computational challenges for both the cognitive modeling and AI communities to wrestle with. We have presented a number of desiderata which correspond well with insights gained in the psychological and neuroscientific literature which we feel are important constraints in developing intelligent systems capable of reasoning about intentions and intentionality.

References

Baldwin, D., and Baird, A. 1999. Action analysis: A gateway to intentional inference. In Rochat, P., ed., *Early Social Cognition: Understanding Others in the First Months of Life*. Mahwah, NJ: Lawrence Erlbaum Associates.

- Bello, P., and Cassimatis, N. 2006a. Developmental accounts of theory-of-mind acquisition: Achieving clarity via computational cognitive modeling. In *Proceedings of the 28th Annual Meeting of the Cognitive Science Society*.
- Bello, P., and Cassimatis, N. 2006b. Understanding other minds: A cognitive modeling approach. In *Proceedings of the 7th International Conference on Cognitive Modeling*.
- Cassimatis, N.; Trafton, J.; Bugajska, M.; and Schultz, A. 2004. Integrating cognition, perception and action through mental simulation in robots. *Journal of Robotics and Autonomous Systems* 49(2):13–23.
- Cassimatis, N. 2002. *Polyscheme: A Cognitive Architecture for Integrating Multiple Representations and Inference Schemes*. Ph.D. Dissertation.
- Cassimatis, N. 2005. Integrating cognitive models based on different computational methods. In *Proceedings of the 27th Annual Meeting of the Cognitive Science Society*.
- Goldman, A. 2006. *Simulating Minds: The Philosophy, Psychology and Neuroscience of Mindreading*. Oxford University Press.
- Jeannerod, M. 2003. The mechanisms of self-recognition in humans. *Behavioral and Brain Research* 142:1–15.
- Meltzoff, A. 2005. Imitation and other minds: the like me hypothesis. In Hurley, S., and Chater, N., eds., *Perspectives on Imitation: from Neuroscience to Social Science*, volume 2. MIT Press. 55–77.
- Murugesan, A., and Cassimatis, N. 2006. A model of syntactic parsing based on domain-general cognitive mechanisms. In *Proceedings of the 28th Annual Meeting of the Cognitive Science Society*.
- Newton, D., and Engquist, G. 1976. The perceptual organization of ongoing behavior. *Journal of Personality and Social Psychology* 12:436–450.
- Nichols, S., and Stich, S. 2003. *Mindreading*. Oxford University Press.
- Rizzolatti, G.; Fadiga, L.; Gallese, V.; and Fogassi, L. 1996. Premotor cortex and the recognition of actions. *Cognitive Brain Research* 3:131–141.
- Ruby, P., and Decety, J. 2001. Effect of subjective perspective taking during the simulation of action: a pet investigation of agency. *Nature Neuroscience* 4:546–550.
- Searle, J. 1983. *Intentionality: An Essay in the Philosophy of Mind*. Cambridge University Press.
- Stroop, J. 1935. Studies of interference in serial verbal reactions. *Journal of Experimental Psychology* 18:622–643.
- Triesman, A., and Gelade, G. 1980. A feature integration theory of attention. *Cognitive Psychology* 12:97–136.
- Woodward, A.; Sommerville, J.; and Guajardo, J. 2001. How infants make sense of intentional action. In Malle, B.; Moses, L.; and Baldwin, D., eds., *Intentions and Intentionality: Foundations of Social Cognition*. Cambridge, MA: MIT Press. 149–169.