Imitation as Analogy

Keith McGreggor
Design Intelligence Laboratory,
School of Interactive Computing, Georgia Institute of Technology,
85 Fifth Street NW, Atlanta, Georgia, USA
bmcgregg@cc.gatech.edu

Abstract

In this paper, I propose a novel cognitive account of imitation as a form of analogical reasoning. I build the case by examining what imitation is, what analogical reasoning is, and the relationship between them. I then illustrate an analogical architecture for imitation, and show how this view of imitation leads to a satisfactory solution of the correspondence problem. Finally, I briefly outline a new system and framework, I•Me•You, for conducting analogical experiments in the imitation context.

Keywords: imitation; cognitive theory; analogical reasoning; theories of mind.

Introduction

> *"Instead of trying to produce a programme to simulate the adult mind, why not rather try to produce one which simulates the child's?" [Turing, 1950]*

Creating a resourceful, resilient software agent with human-level intelligence has been a core goal of artificial intelligence since the field's inception. The problems surrounding this goal are many. We want to understand human cognition, and, in AI, we want to design and engineer human-like cognition into the systems we build. We want our robots and our systems to learn from us, and from each other, efficiently.

We humans do manage to understand each other, despite not being able to fully account for our actions and intentions (or even being aware of them). From our earliest age, we learn much of what we know just by observing our environment and the actions of others in it: to follow one another, to make and use tools, to mirror each other's movements in dance or sports, to speak and laugh and sing, to feel joy or share grief, or to think about things in a similar manner and put ourselves in another's shoes. We imitate each other, and somehow, we learn. Perhaps, if our robots and agents could imitate us the way we imitate each other, they could learn from us like we do.

The problem of creating an imitator is made difficult due to the very different nature of the representations necessary for compiling and associating observed actions into produced movements. Simple perceptual matching of (generally) visual input isn't sufficient to produce the actions we would ascribe to imitation. Despite the tremendous amount of research emphasis given imitation, to date a substantial cognitive account of imitation remains absent.

In this paper, I suggest that imitation is a form of analogical reasoning, and propose a very detailed account of how an agent's observation of action can be shown to cause the production of matching movements. An important distinction between my theory and existing theories is that I attempt to account for precisely how imitation occurs, but not why.

I'll begin with an examination of imitation from a variety of perspectives and provide a review of the existing theoretical literature. I will then briefly discuss analogical reasoning and present my theory of imitation as analogy. I will next describe the I•Me•You system, a framework for conducting imitation and analogical experiments with minimal intelligent agents in embedded within a mini-world. Finally, I will conclude with some thoughts on future steps to take.

What is Imitation?

According to the Oxford English Dictionary, imitation is “the adoption, whether conscious or not, during a learning process, of the behaviour or attitudes of some specific person or model.” [OED] While there are several competing definitions given, it is this definition of imitation that should be understood to be in use throughout this paper. This definition of imitation, however, poses several unique challenges and constraints which need to be explored in detail, if an architecture of imitation is to be constructed.

Architectural Implications

The context of imitation is explicitly that of a learning processing: the agent is the student, and the model, the instructor. There may or may not be an underlying set of motives for the actors of these roles; thus, there need not be a goal other than literal replication. Yet, if this is a learning context, then the agent must have some capacity for retention as well as replication, regardless of the underlying learning mechanism employed. The architectural implication is straightforward: the agent must possess some memory, and that memory must have some organizational structure which affords learning.

Under this definition of imitation, the agent may or may not be aware of its actions. The problem of consciousness or self-awareness is orthogonal to the problem of accounting for imitation. As I explain later, however, the architecture for imitation I propose offers an avenue toward self-awareness.

The agent somehow observes the model’s behavior or attitude, and then adopts some aspect (or aspects) of that model. Here, the definition of imitation poses the most significant challenge. If the agent is adopting a behavior from the model, then the agent must somehow reproduce that behavior (a transformation from observation to configuration of the agent’s body state through action). This suggests that the agent must be somehow embodied in a fashion to permit reproduction of its observations. It is also quite possibly a necessity that the agent be embedded within the same world as the model, or at the very least have its actions subject to observation and comment upon by the world or the model (as this is a learning context: some response must be available to the agent).

A Minimal Architecture of Imitation

The elements of a minimal architecture for imitation within an intelligent agent are:

- An observation layer, through which the agent observes the world from any number of sensory modes;
- A perception layer, which the agent uses to transform its observations into a representational form which may encompass one or more sensory modes;
- A set of internal states, which correspond to an agent’s current configuration (either physical or cognitive);
- A memory, which the agent can use to retain feedback from the world; and
- Some effectors, through which the agent causes actions to occur in the world.

Thus, a minimal architecture for imitation would look something like that given in figure 1:

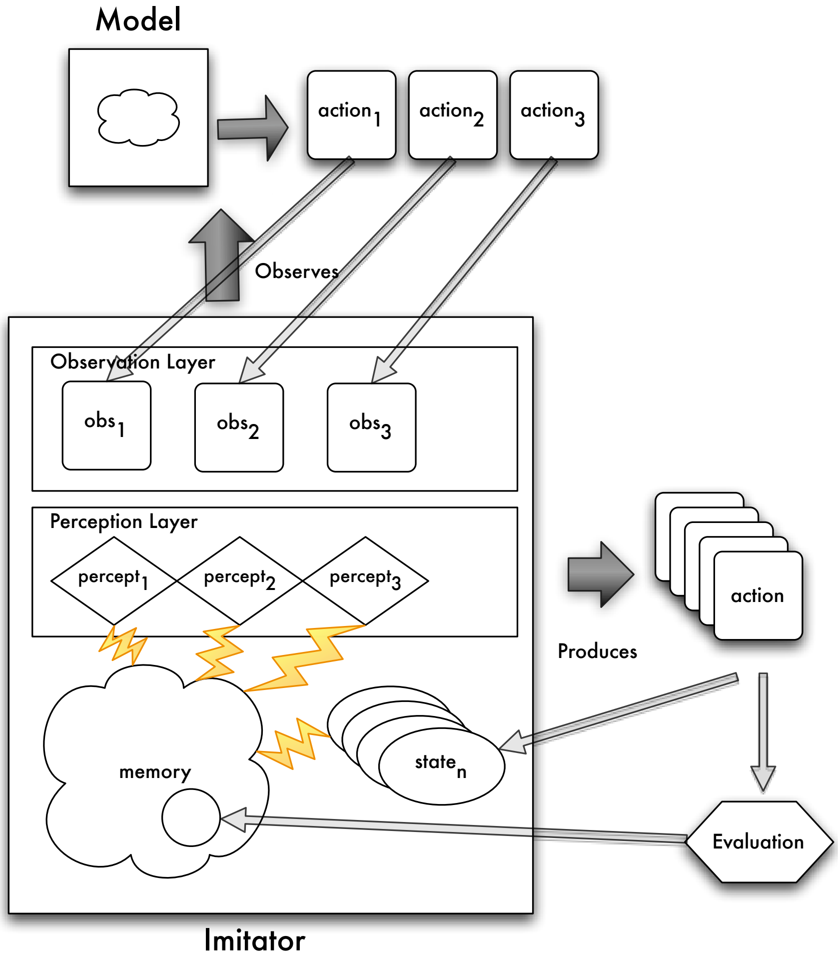

Figure 1. Imitation architecture

The Correspondence Problem

According to Nehaniv and Dautenhahn, to identify social learning, imitation, copying or mimicry presupposes a notion of observable correspondence between two autonomous agents. Judging whether a behavior has been transmitted socially requires an observer to identify a mapping between the model and the imitating agent. [Nehaniv and Dautenhahn, 2002] This is known as the "correspondence problem" which can be more formally stated as follows:

> Given an observed behavior of the model, which from a given starting state leads the model through a sequence (or hierarchy or program) of sub-goals in states, action and/

> or effects, one must find and execute a sequence of actions using ones own (possibly dissimilar) embodiment, which from a corresponding starting state, leads through corresponding sub-goals–in corresponding states, actions and/or effects, while possibly responding to corresponding events. [Nehaniv, 2007]

The correspondence problem and its solution are therefore intricately linked to imitation. Yet, a solution cannot be attempted without first deciding upon a set of criteria or metrics for measuring the similarity between various aspects of the model and the agent, for it is unlikely that the agent and the model will be identical (or even identically enabled or embodied). Nehaniv developed a taxonomy of 24 classes of correspondence problems based upon which aspects of behaviors (the agent's actions, the agent's states, or the effects of the agent's world) and the type of granularity of the changes. [Nehaniv, 2007]

For this proposal, the correspondence problem is significant in exactly one way: it introduces and formalizes the notion of an observer of the resulting activities. Whether the observer is identical to the imitating agent is irrelevant: the existence of an observer provides to the system the necessary architectural imperative to provide a response mechanism to the agent on the success of the imitation, and thus affords learning. The nature of the observer's response then must be such that the information can be rapidly consumed by the agent.

Other Aspects of Imitation

The answer to the question of why an agent would want to imitate another is beyond the scope of this proposal, but I want to address certain aspects of imitation with it in mind.

Imitation and Intrinsic Motivation

Imitation appears to be an intimate part of, yet remains distinct from, intrinsic motivation. Psychologists distinguish between extrinsic motivation (being moved to do something because of some specific rewarding outcome) and intrinsic motivation (being moved to do something because it is inherently enjoyable). Intrinsic motivation leads children to engage in exploration, play, and other behaviors driven by curiosity without explicit reward.

In this research proposal, I suggest that imitation is one of these "other behaviors," driven by some set of low-level internal drives (e.g. seeking novelty), but reinforced by rewarding outcome.

Imitation and Goals

Humans interpret each other's actions in terms of hidden mental states like beliefs, desires, etc. This attribution of intentional mental states to others helps human agents to explain observed, and predict future, behavior, and enables them to influence what their social partners will do [Csibra, 2003].

Imitation is not strictly goal-directed behavior, according to the definition I'm using for this paper, however it may be best to consider it as such to distinguish imitation from mimicry. Dr. Andrea Thomaz, in a personal communication, draws a strong distinction between these two similar notions:

> Imitation is behavior with some end in mind; but mimicry is behavior which results in merely copying the actions of another.

While the specific discovery of goals of the model is not an aim of the proposed architecture for imitation, my proposal accomplishes this readily by adopting a strict teleological read of the "end state of the last action as goal" in the manner of Csibra. At this early stage of the research, however, it remains to be seen whether labeling the goals prove advantageous for any purpose other than having the system explain its reasoning or justify its actions.

Imitation and Social Learning

As Revel and Nadel points out, human infants are equipped to interact dynamically with their physical and social worlds: they are autonomous, novelty expectant, attracted toward movement, and able to couple perception and action [Revel and Nadel, 2007]. Learning by imitation, then, is a powerful and versatile method for acquiring new behaviors.

In human cultures, a wide range of behaviors, from styles of social interaction to tool use, are passed from one generation to another through imitative learning. Tomasello calls this cumulative cultural evolution:

> Basically none of the most complex human artifacts or social practices … were invented once and for all at a single moment by any one individual or group of individuals. Rather, what happened was that some individual or group of individuals first invented a primitive version of the artifact or practice, and then some later user or users made a modification, an "improvement," that others then adopted perhaps without change for many generations, at which point some other individual or group of individuals made another modification, which was then learned and used by others, and so on over historical time in what has sometimes been dubbed "the ratchet effect." [Tomasello, 1999]

Imitation provides not just a mechanism for inheriting behaviors from an agent's social others: it enforces the "ratchet effect" by ensuring that the innovation process continuously find either new, novel ways to improve some goal, or repeat "known" methods for achieving that goal.

Imitation appears a strong supporter of Lamarckian evolution (characteristic behaviors are acquired through generational social interaction and imitation, but they are not (apparently) encoded and passed to future generations genetically) [Rao, et al., 2007], but I disagree somewhat. I believe that there must exist some fundamental repertoire of behaviors in any agent, and through the imitative experience, these behaviors get reinforced. I will revisit this in the details of my proposed architecture for imitation

Unlike trial-and-error based learning methods such as reinforcement learning, imitation allows rapid learning [Rao, et al., 2007]. Imitation it seems can proceed from very few model examples. This implies that the metrics a successfully imitative agent uses to discriminate correspondences must be quite expressive.

Imitation and Simulation

Imitation also appears quite closely related to simulation. According to Gordon, simulation is usually equated with role-taking ("putting oneself in the other's place"). In an intelligent agent which is simulating a model, the agent's output is not imitative behavior but predictions of behavior, such that the agent might find itself "surprised" at the model's eventual pursuit (or lack of pursuit) of a course of action.

Gordon has proposed simulation theory (ST) as a theory "of everyday human psychological competence":

> … of the skills and resources people routinely call on in the anticipation, explanation, and social coordination of behavior. ST holds that we represent the mental activities and processes of others by mentally simulating them, or generating similar activities and processes in ourselves: thus, for example, anticipating another's solution to a theoretical or practical problem by solving the problem ourselves (with adjustments for evident disparities, e.g., in skill level). The basic idea is that if the resources our own brain uses to guide our own behavior can be modified to work as representations of other people, then we have no need to store general information about what makes people tick: We just do the ticking for them." [Gordon, 2008]

Architecturally, an imitative intelligent agent would have to possess at least the capacity of predicting to some degree the outcomes its actions might cause if undertaken, due to the learning context of imitation and the necessity for the incorporation of feedback. If the agent further could allow or inhibit the execution of the imitative action, then perhaps the agent could be said to be simulating the experience of the model's observed actions.

In the proposed architecture, I make a clear distinction between these two notions (imitation and simulation), but show how both exist harmoniously and necessarily in an intelligent agent.

Related Research

Imitation has attracted a tremendous degree of attention and research over the years, from Piaget's theories of the early 1950s to the recent resurgence of activity in imitation research in the arena of social robotics. In this section, I highlight some of the previous work which I felt had bearing on the proposal at hand.

Theories of Imitation and Intrinsic Motivation

According to Piaget, children are intrinsically motivated to seek activities which involve some assimilation and accommodation [Piaget, 1952]. Assimilation is a process whereby children incorporate aspects of the environment into their preexisting cognitive structure, which Piaget called schemata. The implication is that a child's cognitive structure influences her perception of the environment. During accommodation, a child adapts her cognitive structures to better correspond with the perceived environment.

In Piaget's view, assimilation and accommodation constitute learning. When a child receives a percept of the environment which strongly mismatches her existing schemata, assimilation will fail, and the percept will most likely be ignored. When percepts are completely predictable, children generally lose interest in them.

Piaget's work has inspired researchers in artificial intelligence for some time. Many artificial intelligence models make use of internal explicit schema structures under names like frames [Minsky, 1975] or scripts [Schank and Abelson, 1977]. Reinforcement learning depends expressly upon assimilation and accommodation [Taniguchi and Sawargi, 2003].

Singh, Barto, and Chentanez [Singh et al., 2004] presented the results of a computational study of intrinsically motivated reinforcement learning which was aimed at allowing artificial agents to construct and extend hierarchies of reusable skills which are needed for competent autonomy. They note that "an agent having a collection of skills learned through intrinsic reward can learn a wide variety of extrinsically rewarded tasks more easily than an agent lacking these skills."

Existing Cognitive Theories of Imitation

A cognitive account of imitation must include some description of how an agent's observations and perceptions are mapped into actions which can be judged as imitative of the model. Nehaniv [Nehaniv and Dautenhahn, 2007] mentions three specific theories: active intermodal mapping; associative sequence learning; and goal-directed imitation.

Both active intermodal mapping and goal directed imitation theories suggest that the cognitive processes which afford imitation involve some intermediate action representations which are not specifically sensory or related directly to motor control. Active intermodal mapping theory maintains that perceived actions are actively processed to infer the observed model's goals, which are then translated into higher-level representations which are used to produce motor commands. Under goal directed imitation, goals are also extracted from perceived movements, but that goal representation then triggers its most commonly associated motor program, without regard for whether the motor movement matches the movement performed by the observed model. In contrast to these two theories, the Associative Sequence Learning model argues that the perception of action typically prompts the performance of that action directly, without the need for intermediate representation.

Active Intermodal Mapping (AIM)

The Active Intermodal Mapping theory [Meltzoff and Moore, 1997] of imitation introduces three theoretical concepts: organ identification, organ relations, and body babbling. Organ identification is the process by which infants come to identify parts of their body with parts of the bodies of others. Meltzoff and Moore argue that this is the first step in the imitative process. Organ relations refers to the capacity of the infant to parse an observed action into a series of relationships between organs (parts) of the body. The same capacity allows the infant to identify the organ relations of her own body using proprioceptive feedback, and through organ identification, to actively compare the organ relations of the model with her own organ relations. Organ relations provide a common content for the percept of an action to be compared to the action of the perceiver. The third idea, body babbling, refers to the process of learning the relationship between muscle movements and the organ relations which result. Meltzoff and Moore argue that this leads to the formation of a directory of muscle movements and associated organ relations. After such experience-dependent learning, the infant will be able to produce muscle movements leading to specified organ relations.

AIM does not appear to address the correspondence problem, for although the infant can identify the organ relations of her own body, and can identity the organ relations of an observed model's body, the two methods of representations are different. To know when a mismatch has occurred, the representations of the externally perceived organ relations must still be compared to proprioceptively (internally) representation of organ relations. AIM solves this problem by using a higher-level representation to encode organ relations in a percept-general fashion. Externally perceived and proprioceptive organ relations are translated by AIM into this common representational framework, allowing them to be compared directly.

Goal Directed Imitation

The major developmental change in imitation occurs after a few weeks of life when perceived actions are no longer coded as organ relations, but rather as goal-directed actions [Meltzoff and Moore, 1997]. Subsequently, imitation is not of perceived movements, but of inferred goals.

In the goal directed theory of imitation, there is no direct link between the perception and production of body movements: perceived actions are decomposed into a series of aspects (the inferred goal of a movement and the means to achieve it). A child's capacity limitations (perception and recall) imply that only some goal aspects are imitated; movement end-points and the manipulation of objects are more likely to be imitated than either the effector or the movement path. When the goal of the movement has been selected, the movement most commonly associated with this goal will be performed. Thus, there is no special relationship between matching movements. If the observer's most commonly associated movement is the same as that performed by the model, then imitation of the perceived movement will occur. Alternatively, if the movement most commonly associated with the perceived goal is different from the movement of the model, then goal (but not movement) imitation will occur.

Associative Sequence Learning

The Associative Sequence Learning theory [Heyes, 2001] suggests that the correspondence problem is solved through the operation of bidirectional associations between sensory and motor representations of action. Sensory representations are made active when actions are perceived, and they contain information received through the distal senses (vision and audition): they represent what an action "looks like." Motor representations contain motor commands needed to perform the action, and somatosensory (kinesthetic and proprioceptive) information received when the movement is performed: they encode what it "feels like" to perform the action. When a sensory and a motor representation are linked, activation of the sensory representation is propagated to the motor representation. If the sensory and motor components represent the same action, this activation of the motor representation makes imitation possible.

According to the ASL model, while some associations may be innate, the majority of associations are formed through repetitive concurrent activation of sensory and motor representations (contiguity) of the same movement. This experience may consist of concurrent observation and execution of the same movement, leading to a direct association, or it may involve exposure to a common stimulus in conjunction with, on some occasions, observation of the movement, and on other occasions with its execution.

The ASL model does not specify a system which compares sensory and motor representations of movements. To generate imitative behavior, the system does not have to "decide" whether there is a match between the associated sensory and motor representations. It is possible for associative links to be formed between sensory and motor representations of actions that do not match from a third-person perspective, however associations are more likely to be formed between matching rather than non-matching movements due to the environment.

Imitation and Social Robotics

Imitation has garnered quite a bit of research interest from the area of social robotics. Here are synopses of several papers which have specific bearing on this proposal.

Intrinsically Motivated Machines

Kaplan and Oudeyer [Kaplan and Oudeyer, 2007] point out that children seem intrinsically motivated "to manipulate, to explore, to test, to learn and they look for activities and situations that provide such learning opportunities." They outline an architecture of robotic systems which afford the autonomous exploration of their environment not merely to fulfill predefined tasks but "driven by an incentive to search for situations where learning happens efficiently." They are especially interested in why a machine would learn, and what and how it would choose what to learn: a system of artificial curiosity.

Their work relates to imitation in that artificial curiosity can be facilitated by imitation. Festinger's theory of cognitive dissonance [Festinger, 1957] asserted that organisms are motivated to reduce dissonance, that is the incompatibility between internal cognitive structures and the situations currently perceived. A related view was developed later by Kagan [Kagan 1972] that a primary motivation for humans is the reduction of uncertainty in the sense of the "incompatibility between (two or more) cognitive structures, between cognitive structures and experience, or between structures and behavior. This line of inquiry leads directly to various research into human behavior motivation, but resolves into two contrasting points of view: either humans search for some form of optimality as defined by an abstract function or for a particular kind of feeling occurring during challenging situations.

Socially Guided Exploration

In [Breazeal and Thomaz, 2008], the authors suggest that a robot should be able to explore and learn on its own, but also take advantage of guidance from a human partner when available. Their work distinguishes itself from earlier efforts in that most of those efforts are either strongly biased toward guidance (full dependency upon human interaction) or strongly biased toward exploration (where the human's role is to facilitate the reward of behavior). Since the intent is to develop social robots, these robots must be able to operate well at both extremes as well as at any point along what they refer to as the "guidance-exploration spectrum."

With the Socially Guided Exploration system, the robot receives a set of observations, which are then classified into a set of perceptions (with each perception receiving a score in the interval [0,1]). These perceptions are then transformed into beliefs which taken together capture the robot's perceived state. Actions are chosen for performance based upon this state. The system adds to this framework a motivation system based upon notions of novelty and mastery, implemented in a manner which strongly resembles biological homeostasis. These two drives, working in concert, cause the system to explore aspects of its state through acting and observing.

Given this basis, Breazeal and Thomaz construct a reinforcement learner where the tasks the robot performs are sets of actions triggered whenever it is in a particular state (more precisely, when its belief set matches some criteria) and the goals are sets of beliefs which are changed during the task, delimited by which of the belief's perceptions are expected to be a certain value (or in a certain range) and which are expected to be changed by the task (the necessary post condition). The system maintains a set of its known tasks, which it attempts to expand upon by executing one of three competing behaviors (novelty, mastery, or explore). The selection of a learning behavior depends upon both the system's internal motivational state (the interaction between its novelty and mastery drives) and social cues from a human teacher (pointing out either saliency in the environment or suggesting an action for execution).

This work has relevance to imitation and analogy in several ways. The transfer problem is present: how do you map observed suggested behavior into actions? How does the system know to pay attention to the human teacher? In this system, the robot is hardwired to "know" that humans are worth more attention than other objects and that they are expected to provide more information. In a sense, the robot is somewhat imprinted upon the human teacher. The system's confidence weighting (arising from the belief system) for tasks leads to saliency discovery, an important aspect of the discovery of intentionality through imitation.

'Like me': a foundation for social cognition

Meltzoff claims in a more recent paper [Meltzoff, 2007] that the ability to recognize equivalences between ourselves and others is the foundation for social cognition, giving us a way to understand the behaviors we see others perform. Although this paper concerns itself exclusively with human infant studies, these infant studies (and infant imitation study in general) are appropriate for my research, as much (if not most) of human early learning occurs through imitation.

In his paper, Meltzoff reviews a variety of infant studies, carefully noting the observed behavior, and drawing conclusions. In the first study, Meltzoff observed that infants responded positively more frequently when an adult imitated the infant's behavior (for infants, "responded positively" means looked at longer and smiled more at, for they are preverbal). Meltzoff identifies two possible explanations: either the infant responded positively when the adult's actions were temporally contingent with her own, or the infant somehow understood that structure of the adult's movements were similar to her own.

Subsequently, Meltzoff conducted experiments to distinguish between temporal contingency and structural equivalence. He found that infants are able to recognize structural congruence, and in fact, that infants preferred (by looking at and smiling at) adults which acted like them versus adults which acted arbitrarily when they moved. This is the basis of his 'like me' hypothesis.

Meltzoff then looks at whether 'like me' can be extended to perception, and discusses a series of experiments in gaze following using various sorts of blindfolds. The conclusions drawn from these experiments were that infants will generally pick up on referential cues from adults they perceive as 'like me,' but that infants past a certain age will recognize the affordances of objects (like blindfolds which permit vision) and modify their behavior based upon their interpretation of those affordances.

The author then moves into a discussion of a potential cognitive architecture behind the 'like me' phenomena. First, Meltzoff notes that even a very young infant, who has never seen her own face in a mirror, is able to imitate facial expressions (e.g. a tongue thrust) of an adult, and suggests that perception and the production of these actions are "intrinsically intertwined." He further notes, however, that even if the imitation of tongue-thrusting is blocked by a pacifier, prohibiting a resonance-like behavior, the moment the pacifier is removed, the infant will attempt to imitate the action, even if the adult has adopted a neutral facial expression. This, along with other studies, suggest that this imitation ability in infants engages some memory system (requiring a representation of the adult's action).

Meltzoff mentions other experiments that have shown infants will self-correct their actions to produce a more faithful reproduction, and that they show special interest in being imitated themselves (for example, having the adult imitate the infant's self-produced tongue-thrust elicits more positive responses, although the infant had never seen her own face or witnessed her own actions in a mirror). All of these experiments point to an infant possessing the capability of representing the actions of her self and others, and storing and retrieving those representations appropriately. Further, that an infant can perform an imitative action significantly delayed in time from the observation of the action means that the representation of the observation, when retrieved from memory, is somehow mapped into an action that the infant can perform.

These findings of Meltzoff lead him to the conclusion that young infants are able to understand the equivalence between themselves and others before they are verbal or even able to observe their own actions in a mirror. Meltzoff believes this establishes the infant's view of the social world, lets them ascribe intentionality and meaning to others, and forms the basis of empathy.

For my research, this one paper lead me to change my proposed architecture from a single stage association to a multistage case-based reasoner. By taking this step, I believe this enables my architecture to perform simulation, as I will illustrate below.

Teleological and referential understanding

In his paper "Teleological and referential understanding of action in infancy"[Csibra, 2003], Gergely Csibra points out that humans consistently interpret each other's actions by attributing beliefs, desires, and intentions to the other, allowing each of us to explain, understand, predict, and influence each other. Our actions can be interpreted either teleologically (they are about goals, which are the outcome of the action), or referentially (they are highlighting some state of the world). This paper walks through several experiments which show that young infants perform both kinds of interpretation of actions, and that their interpretations become "more sophisticated" with age.

Csibra examines the teleological stance first. The author notes earlier experiments that show very young infants can recognize the end state of a sequence of events, but observes that these experiments did not specifically test the necessity of a particular action. An action is goal-related if it is instrumental to achieving some end. Csibra notes that an action can be intuitively tested as goal-related by considering whether the action would be expected to be performed when unnecessary.

A very interesting (and quite helpful) notion to me was that infants can be tested using computer graphics which greatly simplify (and indeed, remove) the features of agents. So long as the stand-ins for people were apparently self-animate, the infants responded. I view this as particularly good news for my own research, as this suggests that I can focus my attention strictly on the cognitive architecture, and not on aspects of visual (or generally, sensory) recognition.

Csibra set up a series of experiments involving infants watching computer animations of balls performing various actions. Once the infants were habituated to a particular animation, he then would remove a crucial aspect of the world (perhaps a barrier that an animated ball had to hop over), yet left the animation the same (indicating in that experiment that the ball would hop over a now-absent barrier). Csibra found that infants responded positively to (looked at longer) the animations with the now inefficient action than animations where the ball behaved "normally" (without unnecessary hopping).

The interpretation of these experiments was that the infants were adopting a teleological stance in that they were interpreting the actions of the ball for the goal they had been habituated to expect, and when their expectation was not met, they looked longer at the novelty. Csibra states that a teleological representation of an event has three components (behavior, context, and end-state), and can be considered well-formed if and only if the behavior is considered to be an efficient action (a goal-directed action) towards the end state given the context.

Other experiments seemed to confirm that infants care about the efficiency of actions in achieving goals, even if those actions are delayed in time significantly. The three-way association for teleological stance (behavior, context, and end-state) allowed Csibra to stage a variety of experiments to test the importance of each aspect individual. The infant studies confirmed this, even in cases where the end-state was never witnessed by the infants.

The teleological stance, it seems, allows infants to fill in any missing third through inference from the presence of the other two. Csibra notes that the teleological stance is different from an intentional stance (beliefs, desires, intentions) in that the experiments do not imply that the infants are attributing any mental states to the agents within the animation. He notes that the intentional stance would be useful in interpreting actions (or, more generally, performing inference about the teleological stance) but only if the aspects of the intentional stance were derived independently. Infants appear to rely upon behavior cues to identify agents acting in goal-directed ways, and Csibra mentions a study where infants were shown to treat any agent (regardless of features) as goal-directed merely on the basis that the agent should appear to be self-propelled.

Csibra explores the referential stance through similar sets of infant studies, although the studies are dominated by gaze-following problems. One important distinction between this group of experiments and the ones on teleological stance is that the context of the experiment is now dependent upon the establishment of a communicative situation between the model and the infant (either with an adult making sustained eye contact with or a simulation exhibiting some contingent reactivity to the infant). Where the teleological stance looked for evidence of the infants paying attention to efficiency of action, the referential stance experiments look for evidence of directionality (where the infant appears to search for what was being referred to). The referential stance is much more causally linked (agent performs an action, which directs attention to some referent), so the experiments were geared to looking for evidence of novelty when the expected causality was absent.

Both the teleological and referential stances portend significant affect upon my research on imitation architectures. In particular, Csibra has shown a way to make specific inferences in each case. From an imitation point of view, this work leads me to architectural considerations of how goals can be literally inferred, and points out how to categorize the actions, as an aid for analogical retrieval and mapping.

Analogical reasoning

Analogical reasoning is a basic ability thought to be a central component of human cognition [Holyoak, et al., 2001]. It is presumptive reasoning based upon the assumption that if things have some similar attributes, their other attributes will be similar [OED, 2]. Analogical reasoning appears in infants at a very early age, and develops and refines through a human's life. It is apparently ubiquitous, from its use in problem solving, argumentation, and explanation, and is closely related to other basic cognitive processes such as perception, similarity judgment, and metaphor. Given this centrality, the development of cognitive and computational models of analogical reasoning has been given considerable research importance [Kokinov and French, 2002].

The Methodology of Analogical Reasoning

Kokinov and French cleanly delineate the necessary methodology of making analogy. First, you must formulate (or robustly represent) the target problem. Next, you select an appropriate source analog from your previous experiences. You then map the source onto the target according to some similarities (perhaps even those that made you select the source to begin with) and make inferences about the target. You evaluate these inferences, hopefully solving the target problem in doing so. From this, you learn by memorizing either the inferences or the mapping (or both), generalizing or narrowing the information in a way that causes you to reorganize your memory for easier, more efficient retrieval of similar problems in the future.

Case-based reasoning

Case-based reasoning is the process of solving new problems based upon the solutions of similar past problems. Reasoning by reusing past cases is a powerful and frequently applied way to solve problems for humans, a claim is also supported by results from cognitive psychological research [Aamodt and Plaza, 1994]. Roger Schank developed a theory of learning based on retaining experience in a dynamic, evolving memory structure [Schank, 1982], and from this work, Janet Kolodner developed the first case-based reasoning system [Kolodner, 1983]. Case-based reasoning and analogical reasoning are occasionally used synonymously, but this is somewhat imprecise: case-based reasoning is a specialized form of analogical reasoning in which the problems being addressed have the same structure and exist within the same domain.

Generally, case-based reasoning follows a four-step process analogous to that of the methodology of analogical reasoning [Aamodt and Plaza, 1994]:

- Retrieve: Given a target problem, retrieve cases from memory that are relevant to solving it. A case consists of a problem, its solution, and annotations about how the solution was derived.
- Reuse: Map the solution from the previous case to the target problem. This may involve adapting the solution as needed to fit the new situation.
- Revise: Having mapped the previous solution to the target situation, test the new solution in the real world (or a simulation) and, if necessary, revise.
- Retain: After the solution has been successfully adapted to the target problem, store the resulting experience as a new case in memory.

Divisions of analogical reasoning

Analogical reasoning can be decomposed significantly in several ways, depending upon the aspects of the problems to be solved, the manner by which it is initiated, and the context in which it is used.

Analogical reasoning can be divided into two camps: intra-domain and cross-domain. This distinction is made depending upon whether the source and the target problems are of the same domain. Case-based reasoning is an example of intra-domain analogical reasoning.

The cognitive process of analogical reasoning can be initiated through several strategies. If a teacher or some oracle-like system provides an explicit hint that a particular source is to be used, this is a forced analogy. If a source is arbitrarily retrieved from memory and then examined for analogy to the target, this is said to be spontaneous analogy (spontaneous in the sense that nothing about the target was used as an index to retrieve the source). In a related way, constructive analogy systematically transforms a set of sources without apparent applicability to a target problem and evaluating the result. A compound analogy can be formed by repeatedly using other analogy strategies to solve a target problem differentially.

The context in which analogy is used may affect the range of strategies employed in the analogical reasoning process. In problem solving, the basis for a problem's solution may be the solution of an older problem, or it could be the derivation of the previous solution (rather than the solution itself). In argumentation, as an example, an old case is used as an argument that in the new situation, a similar set of solution steps should be applied [Kokinov, n.d.].

Criteria for Similarity

The similarity between the source and the target affects the kinds of mappings available. According to Holyoak and Thagard, similarity, or the judgment of how two situations are analogous, is determined by three constraints: isomorphism, semantic similarity, and pragmatic centrality [Holyoak and Thagard, 1989]. Isomorphism favors mappings which are structural in that they maximize the consistency of correspondence between various elements of the source and the target. Semantic similarity suggests that the source and target are analogous to the degree that the elements and actions are analogous themselves. Pragmatic centrality favors finding correspondence between the most important or salient aspects of the source and the target with respect to the goal of the analogical comparison itself.

Each of these constraints is considered simultaneously, and depending upon the degree to which each is satisfied, the analogical strategy chosen to map from source to target will vary. If the source and target have the same attributes and structure, the analogical mapping can be quite literal. If the source and the target problem have a strong semantic relationship, but no isomorphic commonality, then the analogy may be only superficial [Gentner, 1989].

Structure Mapping and Imitation

Of the criteria for similarity outlined above, isomorphism-structure mapping-holds special significance for imitation. It is very likely the case that while an imitating agent may be similarly embodied and enabled to perform the actions that a model is present, that agent may not be exactly the same in every (physical) regard.

Gentner, in [Gentner, 1983], describes a theoretical framework for analogy built exclusively on isomorphism called structure mapping. In structure mapping, the important aspects are that the rules by which a source is judged similar to a target depend only on syntactic properties of the representation being used, and not upon any domain content. The relations between objects are therfore mapped from the source to the target, and not the attributes of the source. Moreover, the relationships which are mapped are chosen based upon how “higher-order” they appear to be. That is, a structure mapping system would attempt to map the most complex relationship (or relationships) possible between the source and the target. Again, these relationships are entirely syntactic: no special meaning is given to the contents (or values) of the representation except their literal symbolism.

Imitation as Analogy

I propose that imitation is not just a manifestation of analogical reasoning about sensory input, but in particular is an example of two-stage case-based reasoning: one stage for transforming external sensory input into expected internal sensations, and the other for transforming internal sensations into actions.

An analogical architecture of imitation

The architecture I envision for imitation consists of two stages: a stage for transforming external sensory input into expected internal sensations, and another stage for transforming internal sensations into actions. Here is a diagram of the architecture:

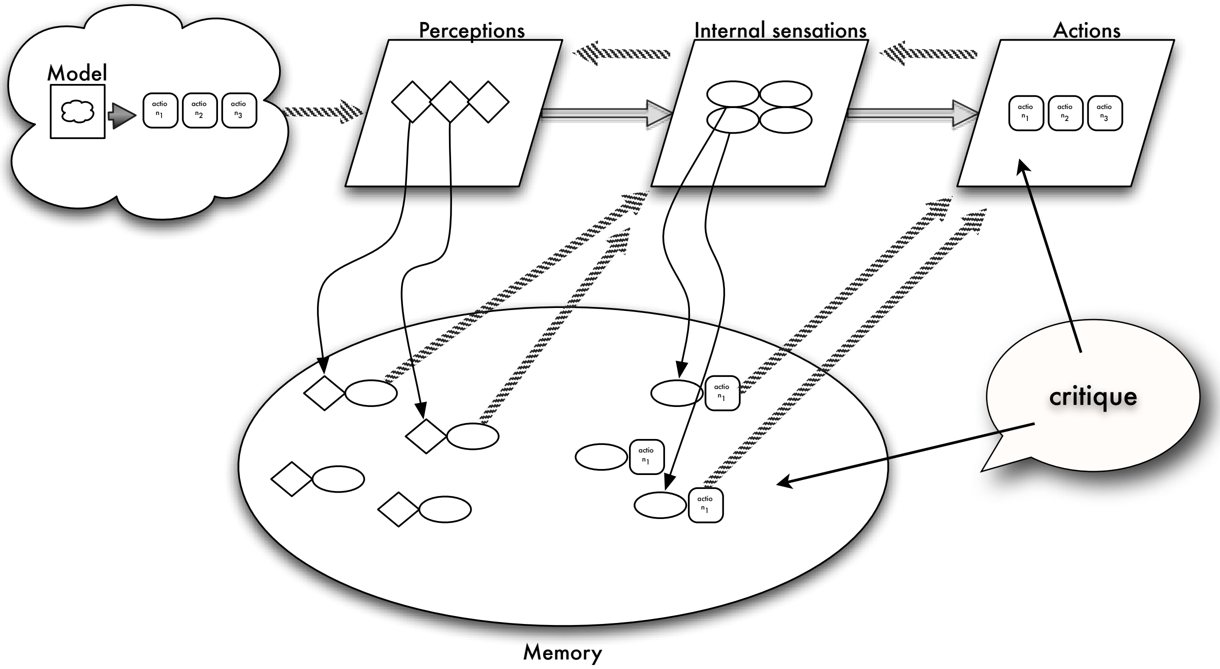

Figure 2. Imitation as Analogy Architecture

To illustrate the functioning of this architecture, I provide the following narrative.

The imitating agent makes some sensory observation (visual, aural, tactile) of the world, in this situation of some model and actions that the model has performed. The agent might perform some processing on this input, to transform the external world observations into some canonical perceptual representation.

Using the perceptual representation, the agent probes its memory for similar previous perceptions, and retrieves cases. These cases contain not only the past perception, but a transcription of the agent's internal sensations (generally prioproceptive) which occurred contingently with those perceptions. Those internal sensations may have been resulted from the agent performing some set of actions in response to the earlier perceptions, but this is not a strict requirement.

The agent maps these earlier internal sensations according to the differences between the past and present perceptions in a structural manner. These newly mapped internal sensations are predictive of how the agent might expect to sense itself.

Next, the agent takes these predicted sensations and, after perhaps combining them with some aspect of its currently experienced internal sensations, probes its memory again for similar past sensations, and retrieves cases. These cases contain the past sensation record as well as a record of what actions the agent performed contingently to cause it to sense those sensations.

The agent maps the retrieved actions according to the differences between the past and present sensations, again in a structural manner. These actions now become the target behavior that the agent will perform. The agent executes the actions and notes how they change its internal sensation or state.

A critic (either the agent itself, or some agency external to the agent) observes the agent's produced actions, compares them with the model's observed actions, and shares this correspondence observation with the agent.

If the agent believes that the way the actions felt correspond well to the retrieved (expected) sensations, the case is reinforced somehow. If the agent feels that they do not correspond well, then a new case is created to capture the experience.

If the external critique (or the reaction of the world) rewards the agent's performance of these actions, the agent reinforces the retrieved external sensory case(s). If the reaction of the world does not reward these actions, the agent may do nothing more, or it might rerun the effort, using the differences noted in the critique as additional information when probing memory.

Comparison to the Methodology of Analogy

Due to the significance of this methodology on my proposed architecture, I'll now dive into each step, commenting on the implications of each.

Representation

As Kokinov and French note, the usual computational system for reasoning analogically does so upon a prescribed or given set of representations of knowledge. Rarely, some systems will form their own representations (or construct re-representations) on unprocessed input [Hofstadter et al., 1995].

In my proposed architecture, the input to the system will be perceptions based upon observations from a variety of sensor readings. The exact nature of the perception's representation is not nearly so important as its ability to discriminate cases within memory. Whether the system eventually manifests a capability to construct its own representations is unclear at this time, but architecturally insignificant.

Retrieval of analogues

Semantic similarity is the most important factor for making the retrieval of past cases efficient and effective. The elements and values of the current percept matter more than the isomorphic resemblance, although the structural similarity can have an impact as well, and will have a more significant impact as the system processes an ever greater variety of cases. Since I plan initially to examine each case in memory for similarity, the simpler the key, the better.

Mapping

Mapping is the core of analogy-making [Kokinov and French, 2002]. While semantic similarity holds sway over retrieval, it is during the mapping segment that isomorphism finds its fullest utility. The structural similarity between the current percept and its retrieved analog are used to map or modify the retrieved sensations to those which should be expected in the first phase of my architecture. Similarly, the transformed (mapped) sensations' similarity to a retrieved analog are mapped onto the retrieved actions for the production of new, prospective actions. I plan to use Gertner's structural mapping process to achieve this.

Transfer and inference

While the transfer and inference at phase one (perception to sensation) is entirely internal, the transfer at phase two has great significance for it creates an expectation of how the produced actions will appear in (and appear to affect) the external world. In practice, this segment likely will be enfolded with the mapping segment.

Evaluation

Classically, the evaluation of the inferences made by analogy tell us whether the analogue was applicable. As I've noted, a critic must exist in the system to be able to make this assessment, based upon its own observation of the actions of the agent as well as its judgment of the correspondence between those actions and those of the model. It is highly likely that at least the initial implementations of this architecture will have the agent act as its own critic. I should note here that the role of the human teacher in the work of Breazeal and Thomaz was to be not just an instructor for the robot in their studies, but would be acting in the critic role for the purposes of this architecture.

Learning

The information coming back from the (possibly external) critic about the correspondence between the agent's actions and those of the model, in addition to the agent's own ability to assess its expected internal sensations with those actually sensed during and after producing the actions, form the basis of powerful feedback for a learning mechanism. I have not yet decided which of several mechanisms will be employed for generalization and memory reorganization.

The Significance of the Stages

Initially, I developed an architecture which created a direct relation between an agent's perceptions and its actions, and this lead to a number of variations which ranged from a strictly hard-wired, reactive interpretation to a more case-based reasoning approach, where previous reactions were preferred. However, as I worked through the architecture, I realized that I was missing an important aspect: I wanted the architecture to afford a simulation capability, that quality of imitation that lets us know what something will feel like just by observing someone doing it. Thus, my proposed architecture has two stages.

By creating a break between the two stages, and providing a means of relating perceptions to internal sensations, this architecture gives the agent a chance to perhaps simulate its chosen actions before (or maybe even without ever) doing them. This is extremely significant in that this one mechanism provides a way for the agent to project its expected internal sensations onto the model - the recognition of "others as same." Again, this arises not due to any particularly special mechanism, but due to the concepts being associated as a consequence of case-based reasoning over perception/sensation tuples.

Imitation and the varieties of analogy

My architecture for imitation as analogy provides a view of analogical reasoning that fuses both intra- and cross-domain reasoning. Taken independently, each stage is performing a very precise intra-domain analogical reasoning process (case-based reasoning), but the system as a whole is very much a cross-domain reasoner. The two stages shift from a representation of perceptions to a concept of actions by first passing through a transformation into (and out of) internal sensations.

The notion of strategies are bound exclusively to the retrieval of source analogs from memory at each stage. Due to the separation of stages, different strategies can be employed for each as situations warrant. For example, it might be that the number of internal sensation cases is significantly less than the number of perception cases. In that situation, an exhaustive, detailed exploration of each case for similarity might be afforded, whether issues of efficiency might preclude such an exhaustive search for similar perceptions. This is analogous to the efficiency rationale used by Marr in the development of his visual processing system [Marr, 1980].

Imitation as analogy and compound analogy

Newly minted agents may lack the experience of cases necessary to sufficiently imitate as desired. For these agents, the notions of compound analogical reasoning and partial solution come into play. The separation of the architecture into the two stages affords an iteration to occur at the first stage (via simulation and internal feedback) before any actual action is produced. Another instance of partial solution would occur if the agent's produced actions only partially correspondence to that of the model. An appropriate feedback mechanism could force the agent into a series of repetitive actions which converge upon a closer facsimile.

Intrinsic analogical reasoning

This cognitive reworking of imitation addresses and complements the intrinsic motivation research. It provides an additional behavior (analogical reasoning) for an intelligent to perform while exploring its world for novelty. The architecture also affords a mechanism for the agent to take either a teleological or referential stance and discover goals (in this case, the direct association of a perception with an action of another, through a single, salient internal sensation set).

Hardwired analogies

The choice of a case-based system may seem odd given that we must populate such a system with cases before anything meaningful can happen. I believe that by doing just that - populating the memory of the agent with readily retrieved actions - this architecture creates a way of accounting for intrinsic, hard-wired behaviors, especially those devoted to survival strategies (eat and flight). I believe human infants come prewired with certain cases, readily retrievable, and naturally or culturally evolved. We have a prebuilt repertoire of actions, and through an architectural mechanism like the one I've outlined above, we refine their use.

The continuum of reasoning

In our research group, we look at analogical reasoning as a continuum from visual reasoning (thinking in pictures) right on up through symbolic reasoning. By casting imitation as analogy, I now add a new, lowest rung on that ladder - sensory/physical reasoning - and we might begin to think of the same fundamental analogical mechanism empowering all three. In my mind, I'm now calling this "thinking in senses" or "sensical thinking."

I • Me • You, a cognitive framework for imitation research

This research proposal suggests the creation of a framework for investigation of imitation as pure analogical reasoning. The framework, named I•Me•You, creates an artificial world populated with intelligent agents, each of which is equipped with sensory and motor capabilities.

I've just begun to work on the first iteration of this framework. I've begun writing it in Java, but I've already tinkered with aspects of the framework in Javascript, under the auspices of another project I conducted with Dr. Thomaz.

Ollie

Ollie (named for Oliver Johnston, one of Disney's famous "nine old men" animators [Thomas and Johnston, 1981]) is a simple web-borne agent for conducting limited experiments in imitation. Ollie's world is a 1-D or 2-D set of pixels, over which Ollie scans his visual system (currently, exactly one pixel in size). Ollie can either scan arbitrarily, or can be directed to look for novelty.

The web page in which Ollie lives allows the user to interact directly with Ollie's world by passing the mouse over various pixels, changing their color in the process. Alternatively, the user can turn on one of several experiments which will introduce objects that move into Ollie's world.

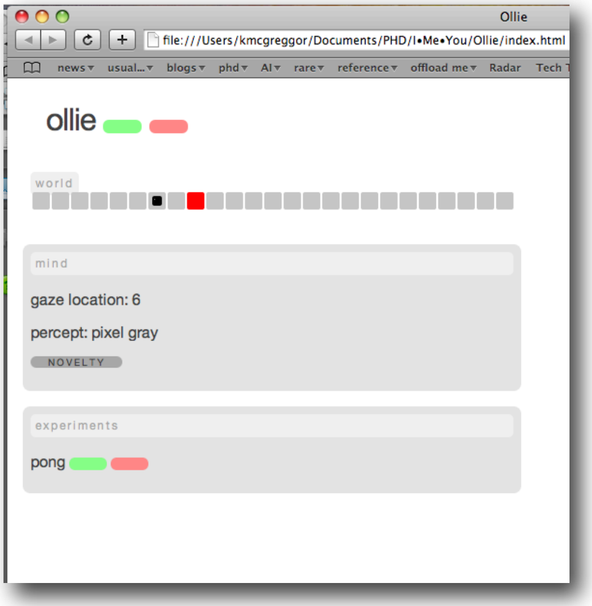

gaze location: 6

percept: pixel gray

Figure 3. Ollie's interface

Even its coding is in a primitive state, Ollie can be demonstrated to track objects (or the user's mouse) within its world. The development of Ollie has just begun, but the intention is to replicate the Csibra experiments in Ollie's world, and give Ollie the foundations necessary to explain a teleological interpretation of its world. Other experiments, with interaction from the user, can be designed to illustrate Ollie adopting a referential stance, or perhaps learning certain tasks ala the Breazeal/Thomaz architecture [Breazeal and Thomaz, 2008].

Conclusions and Future Work

The aim of this paper has been to show that the role analogy plays in imitation has significant strengths as a cognitive account that are not to be found in existing theories. In addition, while the examination of previous research into imitation should not be construed as exhaustive by any metric, it does appear that no one has proposed creating such an account of imitation in either cognitive science, classical psychology, or artificial intelligence. This approach to understanding and recreating imitation seems, for the moment, to be a unique and novel cognitive and computational account.

This is just the beginning of what I hope will be a long and productive cycle. I've quite the stack of possibly related research to analyze, especially in the areas of social robotics, cultural cognition and neuroscience areas. I'm particularly interested in looking at the latest developments concerning mirror neurons, to see if there may be some avenue for deepening the continuum of analogy.

I will also be ramping up my coding efforts. The next step is to fully flesh out the architecture of I•Me•You in Java, paying special attention to any inadvertent bias that could creep into the coding (either bias in the control structures or in the representations of perceptions, internal sensations and states, or memory). While I believe that a two-stage process for transforming perceptions into actions is sufficient, I will be looking for additional evidence to support that belief or rationales for adding additional stages.

Once initial coding is complete, I'm going to create a mini-world in which I hope to reproduce some of the classic human infant imitation studies, especially those experiments of Csibra mentioned earlier. Orthogonal to this, I will be refining the cognitive architecture, to ensure that the only enabling aspects are analogical. I hope one day to install I•Me•You as the control system of an actual social robot, and conduct real-world imitation studies.

Acknowledgements

I would like to thank Ashok Goel, Andrea Thomaz, and the members of the Design Intelligence Lab for their many helpful comments and conversations.

References

[Aamodt and Plaza, 1994] Aamodt, A., and Plaza, E. (1994). Case-Based Reasoning: Foundational Issues, Methodological Variations, and System Approaches. *AI Communications. IOS Press, 7(1)*, 39-59.

[Brezeal and Thomaz, 2008] Breazeal, C., and Thomaz, A.L. (2008). Learning from human teachers with socially guided exploration. *Proceedings of the 2008 IEEE International Conference on Robotics and Automation (ICRA-08)*, 3539-3544. Pasadena, CA.

[Csibra, 2003] Csibra, G. (2003). Teleological and referential understanding of action in infancy. *Philosophical Transactions of the Royal Society B: Biological Sciences, 358*, 447-458.

[Festinger, 1957] Festinger, L. (1957). *A theory of cognitive dissonance*. Evanston, IL: Row, Peterson.

[Gentner, 1983] Gentner, D. (1983). Structure-Mapping: A Theoretical Framework for Analogy. *Cognitive Science, 7(2)*, 155-170

[Gentner, 1989] Gentner, D. (1989). The Mechanisms of Analogical Learning. In Vosniadou S., and Ortony A. (eds.) *Similarity and Analogical Reasoning*. New York, NY: Cambridge Univ. Press.

[Gordon, 2008] Gordon, R.M. (2008). Folk Psychology as Mental Simulation. *The Stanford Encyclopedia of Philosophy (Winter 2008 Edition)*, Edward N. Zalta (ed.), forthcoming. <http://plato.stanford.edu/archives/win2008/entries/folkpsych-simulation/>

[Heyes, 2001] Heyes, C. (2001). Causes and consequences of imitation. *Trends in Cognitive Science, 2001; 5(6)*, 253-261.

[Hofstadter et al., 1995] Hofstadter, D. and the Fluid Analogies Research Group (1995). *Fluid Concepts and Creative Analogies: Computer Models of the Fundamental Mechanisms of Thought.* New York, NY: Basic Books.

[Holyoak, et al., 2001] Holyoak, K., Gentner, D., Kokinov, B. (2001). The Place of Analogy in Cognition. In Holyoak, K., Gentner, D., Kokinov, B. (eds.) *The Analogical Mind: Perspectives from Cognitive Science*. Cambridge, MA: MIT Press.

[Holyoak and Thagard, 1989] Holyoak K. and Thagard P. (1989). Analogical Mapping by Constraint Satisfaction. *Cognitive Science, 13,* 295-355.

[Kagan, 1972] Kagan, J. (1972). Motives and development. *Journal of Personality and Social Psychology 22*, 51–66.

[Kaplan and Oudeyer, 2007] Kaplan, F. and Oudeyer, P-Y. (2007, to appear). Intrinsically Motivated Machines. In M. Lungarella and Iida, F, Bongard, J,. and Pfeifer, R. (Eds.): *50 Years of AI, Festschrift, LNAI 4850*, 304–315.

[Kokinov, n.d.] Kokinov, B. (n.d.) *Analogy-Making: Psychological Data and Computational Models.* Retrieved on December 5, 2008 from Web Site: <http://www.nbu.bg/cogs/personal/kokinov/COG501/Analogy-making.html>

[Kokinov and French, 2002] Kokinov, B., and French, R. (2002). Computational Models of Analogy-Making. In Lynn Nadel (ed.) *Encyclopedia of Cognitive Science*. Macmillan.

[Kolodner, 1983] Kolodner, J. (1983). Maintaining organization in a dynamic long-term memory. *Cognitive Science, 7*, 243-280.

[Marr 1980] Marr, D. (1980). Visual information processing: the structure and creation of visual representations. *Philosophical Transactions of the Royal Society B: Biological Sciences 290,* 199-218.

[Meltzoff and Moore, 1997] Meltzoff, A.N., and Moore, M.K. (1997). Explaining facial imitation: A theoretical model. *Early Development and Parenting, 6,* 179–192.

[Meltzoff, 2007] Meltzoff, A.N. (2007). 'Like Me': a foundation for social cognition. *Developmental Science, 10 (1),* 126-134.

[Minsky, 1975] Minsky, M. (1975) A framework for representing knowledge. In: Winston, P.H. (ed.) *The psychology of computer vision*, 211–277. New York: McGraw Hill.

[Nehaniv and Dautenhahn, 2002] Nehaniv, C. L., and Dautenhahn, K. (2002) The Correspondence Problem. *Imitation in Animals and Artifacts, MIT Press, 2002.*

[Nehaniv, 2007] Nehaniv, C.L. (2007) Nine Billion Correspondence Problems. In Nehaniv, C. L. and Dautenhahn, K. (Eds.), *Imitation and Social Learning in Robots, Humans and Animals: Behavioural, Social and Communicative Dimensions*. Cambridge University Press.

[Nehaniv and Dautenhahn, 2007] Nehaniv, C.L. and Dautenhahn, K. (2007) *Imitation and Social Learning in Robots, Humans and Animals: Behavioural, Social and Communicative Dimensions.* Cambridge University Press.

[OED] "imitation, n.1c" The Oxford English Dictionary. 2nd ed. 1989. *OED Online*. Oxford University Press. 4 Apr. 2000 <http://dictionary.oed.com/cgi/entry/50112379>.

[OED,2] "analogy, 7b" The Oxford English Dictionary. 2nd ed. 1989. *OED Online*. Oxford University Press. 4 Apr. 2000 <http://dictionary.oed.com/cgi/entry/50007888>.

[Piaget, 1952] Piaget, J.: The origins of intelligence in children. Norton, New York (1952)

[Rao, et al., 2007] Rao, R.P.N, Shon, A.P., and Meltzoff, A.N. (2007). A Bayesian model of imitation in infants and robots. In Nehaniv, C. L. and Dautenhahn, K. (Eds.), *Imitation and Social Learning in Robots, Humans and Animals: Behavioural, Social and Communicative Dimensions*. Cambridge University Press.

[Revel and Nadel, 2007] Revel, A., and Nadel, J. (2007). How to Build an Imitator. In Nehaniv, C. L. and Dautenhahn, K. (Eds.), *Imitation and Social Learning in Robots, Humans and Animals: Behavioural, Social and Communicative Dimensions*. Cambridge University Press.

[Schank and Abelson, 1977] Schank, R., Abelson, R. (1977). *Scripts, plans, goals and understanding: An inquiry into human knowledge structures.* Hillsdale, NJ: Lawrence Erlbaum Associates.

[Schank, 1982] Schank, R. (1982). *Dynamic memory; a theory of reminding and learning in computers and people.* Cambridge University Press.

[Singh et al., 2004] Singh, S., Barto, A. G., and Chentanez, N. (2004). Intrinsically motivated reinforcement learning. In *Advances in Neural Information Processing 18*.

[Taniguchi and Sawargi, 2003] Taniguchi, T., and Sawaragi, T. (2003). Assimilation and Accommodation for Self-organizational Learning of Autonomous Robots: Proposal of Dual-Schemata Model. *IEEE International symposium on CIRA 2003*, 277-282.

[Thomas and Johnston, 1981] Thomas, F., and Johnston, O. (1981) *Disney Animation: The Illusion of Life.* Abbeville Press.

[Tomasello, 1999] Tomasello, M. (1999). *The Cultural Origins of Human Cognition.* Cambridge, MA: Harvard University Press.

[Turing, 1950] Turing, A. (1950). Computing machinery and intelligence. *Mind 59,* 433–460.