

Stages of Cognitive Development in Uncertain-Logic-Based AI Systems

Ben Goertzel¹, Stephan Vladimir Bugaj²

¹Novamente, LLC and Virginia Tech, ²AGI Research Institute

ben@goertzel.org, stephan@bugaj.com

Abstract

A novel theory of stages in cognitive development is presented, loosely corresponding to Piagetan theory but specifically oriented toward AI systems centered on uncertain inference components. Four stages are articulated (infantile, concrete, formal and reflexive), and are characterized both in terms of external cognitive achievements (a la Piaget) and in terms of internal inference control dynamics. The theory is illustrated via the analysis of specific problem solving tasks corresponding to the different stages. The Novamente AI Engine, with its Probabilistic Logic Networks uncertain inference component and its embodiment in the AGI-SIM simulation world, is used as an example throughout.

I. INTRODUCTION

Contemporary cognitive science contains essentially no theory of “AI developmental psychology” – a lack which is frustrating from the perspective of AI scientists concerned with understanding, designing and controlling the cognitive development of generally intelligent AI systems. There is of course an extensive science of human developmental psychology, and so it is a natural research program to take the chief ideas from the former and inasmuch as possible port them to the AI domain. However this is not an entirely simple matter both because of the differences between humans and AI’s and because of the unsettled nature of contemporary developmental psychology theory. The present paper describes some work that we have done in this direction, as part of a longer-term project to develop a systematic theory of AI cognitive development.

The ghost of Jean Piaget hangs over modern developmental psychology in a yet unresolved way. Piaget’s theories provide a cogent overarching perspective on human cognitive development, coordinating broad theoretical ideas and diverse experimental results into a unified whole. Modern experimental work has shown Piaget’s ideas to be often oversimplified and incorrect.

However, what has replaced the Piagetan understanding is not an alternative unified and coherent theory, but a variety of microtheories addressing particular aspects of cognitive development. For this reason a number of contemporary theorists taking a computer science [1] or dynamical systems [2-4] approach to developmental psychology have chosen to adopt the Piagetan framework in spite of its demonstrated shortcomings, both because of its conceptual strengths and for lack of a coherent, more rigorously grounded alternative.

The work described here involves the construction of a theory of cognitive development inspired conceptually by Piaget’s work, but specifically applicable to AI systems that rely on uncertain logical inference as a primary or highly significant component. Piaget describes a series of stages of cognitive development, each corresponding to a certain level of sophistication in terms of the types of reasoning a child can carry out. We describe a related series of stages, each corresponding not only to a level of sophistication in terms of demonstrated problem-solving ability, but also to a level of internal sophistication in terms of inference control mechanisms within AI software implementations.

This work was inspired by our ongoing research involving the Novamente AI Engine [5-7], a complex integrative software system aimed at achieving advanced Artificial General Intelligence (AGI) [8]. The Novamente system has been integrated with AGI-SIM, a 3D simulation world powered by the CrystalSpace game engine used in the Crystal Cassie embodiment of the SNePs AGI system [9]. The table below shows each of our proposed developmental stages, with examples drawn from our ongoing research with the Novamente system:

Stage	Example
Infantile	Object Permanence
Concrete	Conservation of Number, Theory of Mind
Formal	Systematic Experimentation
Reflexive	Correction of Inference Bias

II. PIAGET’S APPROACH TO COGNITIVE DEVELOPMENT

Jean Piaget, in his classic studies of human developmental psychology [10-15], conceived of child development in four

stages, each roughly identified with an age group: infantile, preoperational, concrete operational, and formal.

--*Infantile*: In this stage a mind develops basic world-exploration driven by instinctive actions. Reward-driven reinforcement of actions learned by imitation, simple associations between words and objects, actions and images, and the basic notions of time, space, and causality are developed. The most simple, practical ideas and strategies for action are learned.

--*Preoperational*: At this stage we see the formation of mental representations, mostly poorly organized and unabstracted, building mainly on intuitive rather than logical thinking. Word-object and image-object associations become systematic rather than occasional. Simple syntax is mastered, including an understanding of subject-argument relationships. One of the crucial learning achievements here is "object permanence"--infants learn that objects persist even when not observed. However, a number of cognitive failings persist with respect to reasoning about logical operations, and abstracting the effects of intuitive actions to an abstract theory of operations.

--*Concrete*: More abstract logical thought is applied to the physical world at this stage. Among the feats achieved here are: reversibility--the ability to undo steps already done; conservation--understanding that properties can persist in spite of appearances; theory of mind--an understanding of the distinction between what I know and what others know. (If I cover my eyes, can you still see me?) Complex concrete operations, such as putting items in height order, are easily achievable. Classification becomes more sophisticated, yet the mind still cannot master purely logical operations based on abstract logical representations of the observational world.

--*Formal*: Abstract deductive reasoning, the process of forming, then testing hypotheses, and systematically reevaluating and refining solutions, develops at this stage, as does the ability to reason about purely abstract concepts without reference to concrete physical objects. This is adult human-level intelligence. Note that the capability for formal operations is intrinsic in the PTL component of Novamente, but in-principle capability is not the same as pragmatic, grounded, controllable capability.

Despite the influence and power of Piaget's theory, it has received much valid criticism. Very early on, Vygotsky [16, 17] disagreed with Piaget's explanation of his stages as inherent and developed by the child's own activities, and Piaget's prescription of good parenting as not interfering with a child's unfettered exploration of the world. Much of the analysis of Piaget's stages as being asocially grounded start with Vygotsky's assertion that children function in a world surrounded by adults who provide a cultural context, offering ongoing assistance, critique, and ultimately validation of the child's developmental activities.

Vygotsky also was an early critic with respect to the idea that cognitive development is continuous, and continues beyond Piaget's formal stage. Gagne [18] also believes in continuity, and that learning of prerequisite skills made the learning of subsequent skills easier and faster without regard

to Piagetan stage formalisms. Subsequent researchers have argued that Piaget has merely constructed ad hoc descriptions of the sequential development of behaviour [19-22]. We agree that learning is a continuous process, and our notion of stages is more statistically constructed than rigidly quantized.

Critique of Piaget's notion of transitional "half stages" is also relevant to a more comprehensive hierarchical view of development. Some have proposed that Piaget's half stages are actually stages [23]. As Commons and Pekker [22] point out: "the definition of a stage that was being used by Piaget was based on analyzing behaviors and attempting to impose different structures on them. There is no underlying logical or mathematical definition to help in this process..." Their Hierarchical Complexity development model uses task achievement rather than ad hoc stage definition as the basis for constructing relationships between phases of developmental ability--an approach which we find useful, though our approach is different in that we define stages in terms of specific underlying cognitive mechanisms.

Another critique of Piaget is that one individual's performance is often at different ability stages depending on the specific task (for example [24]). Piaget responded to early critiques along these lines by calling the phenomenon "horizontal décalage," but neither he nor his successors [25,26] have modified his theory to explain (rather than merely describe) it. Similarly to Thelen and Smith [2], we observe that the abilities encapsulated in the definition of a certain stage emerge gradually during the previous stage--so that the onset of a given stage represents the mastery of a cognitive skill that was previously present only in certain contexts.

Piaget also had difficulty accepting the idea of a preheuristic stage, early in the infantile period, in which simple trial-and-error learning occurs without significant heuristic guidance [27], a stage which we suspect exists and allows formulation of heuristics by aggregation of learning from preheuristic pattern mining. Coupled with his belief that a mind's innate abilities at birth are extremely limited, there is a troublingly unexplained transition from inability to ability in his model.

Finally, another limiting aspect of Piaget's model is that it did not recognize any stages beyond formal operations, and included no provisions for exploring this possibility. A number of researchers [25,28-31] have described one or more postformal stages. Commons and colleagues have also proposed a task-based model which provides a framework for explaining stage discrepancies across tasks and for generating new stages based on classification of observed logical behaviors. [32] promotes a statistical conception of stage, which provides a good bridge between task-based and stage-based models of development, as statistical modeling allows for stages to be roughly defined and analyzed based on collections of task behaviors.

[29] postulates the existence of a postformal stage by observing *elevated levels of abstraction* which, they argue, are not manifested in formal thought. [33] observes a postformal stage when subjects become capable of analyzing

and coordinating complex logical systems with each other, creating metatheoretical supersystems. In our model, with the reflexive stage of development, we expand this definition of metasystemic thinking to include the ability to consciously refine one's own mental states and formalisms of thinking. Such self-reflexive refinement is necessary for learning which would allow a mind to analytically devise entirely new structures and methodologies for both formal and postformal thinking.

III. THE UNCERTAIN INFERENCE PARADIGM

Piaget's developmental stages are very general, referring to overall types of learning, not specific mechanisms or methods. This focus was natural since the context of his work was *human* developmental psychology, and neuroscience has not yet progressed to the point of understanding the neural mechanisms underlying any sort of inference. But if one is studying developmental psychology in an AI context where one knows something about the internal mechanisms of the AI system under consideration, then one can work with a more specific model of learning. Our focus here is on AI systems whose operations contain uncertain inference as a central component.

An uncertain inference system, as we consider it here, consists of four components:

- a content representation scheme
- an uncertainty representation scheme
- a set of inference rules
- a set of inference control schemata

Examples of content representation schemes are predicate logic and term logic [34]. Examples of uncertainty representation schemes are fuzzy logic [35], imprecise probability theory [36,37], Dempster-Shafer theory [37,38], Bayesian probability theory [39], NARS [40], and the Probabilistic Logic Networks (PLN) representation used in Novamente [41].

Many, but not all, approaches to uncertain inference involve only a limited, weak set of inference rules (e.g. not dealing with complex quantified expressions) Both NARS and PLN contain uncertain inference rules that apply to logical constructs of arbitrary complexity.

The subtlest part of uncertain inference is inference control: the choice of which inferences to do, in what order. Inference control is the primary area in which human inference currently exceeds automated inference. Humans are not very efficient or accurate at carrying out inference rules, with or without uncertainty, but we are very good at determining which inferences to do and in what order, in any given context. The lack of effective, context-sensitive inference control heuristics is why the general ability of current automated theorem provers is considerably weaker than that of a mediocre university mathematics major [42].

IV. NOVAMENTE AND PROBABILISTIC LOGIC NETWORKS

Novamente's knowledge representation consists of weighted labeled, generalized hypergraphs. Patterns embodying knowledge emerge from applying various learning and reasoning algorithms to these hypergraphs.

A hypergraph is an abstract mathematical structure, which consists of objects called Vertices and objects called Edges, which connect the Vertices [43]. In Novamente we have adopted the terminology of using *Node/Vertex* to refer to the elements of the hypergraph that are concretely implemented in a Novamente system's memory, and *Link/Edge* to refer to elements of hypergraphs that are used to model Novamente systems and represent patterns that emerge in the concretely implemented hypergraph. We use the term *Atom* to refer to Nodes and Links inclusively. A hypergraph differs from a graph in that it allows Edges to connect more than two Vertices. Novamente hypergraphs extend ordinary hypergraphs to contain additional features, such as Edges that point to Edges instead of Vertices, and Vertices that represent complete sub-hypergraphs.

A "weighted, labeled hypergraph" is a hypergraph whose Atoms all have associated annotations called *labels*, and one or more numbers that are generically called *weights*. The label associated with an Atom might be interpreted as telling you what *type* of entity it is (a metalogical knowledge annotation). An example of a weight attached to an Atom is a number representing a probability, or a number representing how important the Atom is to the system.

In the framework introduced in the previous section, Novamente's content representation is a "labeled generalized hypergraph with weights representing the attention paid to hypergraph components via learning and reasoning algorithms" and the uncertainty representation consists of some additional weights attached to the Nodes and Links of the hypergraph, representing probability values and related quantities such as "weight of evidence."

Novamente's knowledge representation includes various types of Nodes, including ConceptNodes and SchemaNodes. SchemaNodes embody cognitive, perceptual or motoric procedures, and are represented as mathematical objects using arithmetic, logical and combinatory operators to combine elementary data types and Novamente Nodes and Links. It also includes a number of other node types including PredicateNodes (SchemaNodes that produce truth values as their outputs) and Nodes representing particular kinds of concrete information, such as NumberNodes, WordNodes, PolygonNodes, etc. An extensive list is given in [6].

Novamente also contains a variety of Link types, including some that represent logical relationships, such as ExtensionalInheritanceLink (ExtInhLink: an edge which indicates that the source Atom is a special case of the target), ExtensionalSimilarityLink (ExtSimLink: which indicates that one Atom is similar to another), and ExecutionLink (a ternary edge, which joins {S,B,C} when S is a SchemaNode and the result from applying S to B is C). Thus, a Novamente knowledge network is a hypergraph whose Nodes represent ideas or procedures, and whose Links represent relationships of specialization, similarity or

transformation among ideas and/or procedures.

ExtInh and ExtSim Links come with probabilistic weights indicating the extent of the relationship they denote (e.g. the ExtSimLink joining the “cat” ConceptNode to the “dog” ConceptNode gets a higher probability weight than the one joining the “cat” ConceptNode to the “washing machine” ConceptNode). The mathematics of transformations involving these probabilistic weights becomes quite involved--particularly when one introduces SchemaNodes corresponding to abstract mathematical operations. SchemaNodes enable Novamente hypergraphs to have the complete mathematical power of standard logical formalisms like predicate calculus, but with the added advantage of a natural representation of uncertainty in terms of probabilities, as well as a neurostructurally motivated model of complex knowledge as dynamical networks.

Novamente contains a probabilistic reasoning engine called Probabilistic Logic Networks (PLN) which exists specifically to carry out reasoning on these relationships, and will be described in a forthcoming publication [8]. The mathematics of PLN contains many subtleties, and there are relations to prior approaches to uncertain inference including NARS [40] and Walley’s theory of interval probabilities [44]. The current implementation of PLN within the Novamente software has been tested on various examples of mathematical and commonsense inference.

A simple example of a PLN uncertain inference rule is the probabilistic deduction rule, which takes the form

$A \rightarrow B$
 $B \rightarrow C$
 $\vdash$
 $A \rightarrow C$

(where e.g. $A \rightarrow B$ is a shorthand for the ExtInhLink from A to B), whose uncertain truth value formula has as one component the formula

$$s_{AC} = s_{AB} s_{BC} + (1-s_{AB}) (s_{C--} s_B s_{BC}) / (1-s_B)$$

(where e.g. s_{AC} and s_B refer to the probability values attached to $A \rightarrow C$ and B respectively). PLN attaches to each node and link a “weight of evidence” value in addition to a probability, but the deduction formula for weight of evidence is more complex and will not be given here.

Inference control in Novamente takes several forms:

1. Standard forward-chaining and backward-chaining inference heuristics (see e.g. [45])
2. A reinforcement learning mechanism that allows inference rules to be chosen based on experience. Probabilities are tabulated regarding which inference rules have been useful in the past in which contexts, and these are subsequently used to bias the choices of inference rules during forward or backward chaining inference
3. Application of PLN inference to the probabilities used in the reinforcement learning mechanism--this

enables generalization, abstraction and analogy to be used in guessing which inference rules may be most useful in a given context

These different approaches to inference control enable increasingly complex inferences, and involve increasing amounts of processor-time utilization and overall cognitive complexity. They may also be interpreted as corresponding to loosely Piagetan stages of cognitive development.

V. DEFINING DEVELOPMENTAL STAGES IN TERMS OF INFERENCE CONTROL

Inspired by Piaget’s general ideas, later critiques, and the structure of inference control in Novamente, we have created a novel theory of cognitive developmental stages, defined in terms of the control of uncertain inference trajectories. The stages are defined as follows:

--*Infantile*: Able to recognize patterns in and conduct inferences about the world, but only using simplistic hard-wired (not experientially learned) inference control schema, along with pre-heuristic pattern mining of experiential data.

--*Concrete*: Able to carry out more complex chains of reasoning regarding the world, via using inference control schemata that adapt behavior based on experience (reasoning about a given case in a manner similar to prior cases).

--*Formal*: Able to carry out arbitrarily complex inferences (constrained only by computational resources) via including inference control as an explicit subject of abstract learning.

--*Reflexive*: Capable of self-modification of internal structures. (In the case of a Novamente, this process is very direct and thorough.)

Here Piaget’s preoperational phase appears as transitional between the infantile and concrete operational phases. We suspect this approach to cognitive modeling may have general value beyond Novamente, but we will address a more generalized developmental theory in future writings. We have designed specific Novamente / AGI-SIM learning tasks based on all the key Piagetan themes. Currently our concrete work is near the beginning of this list, at Piaget’s infantile stage.

The semantics of our stages is similar but not identical to Piaget’s. Our stages are defined via internal cognitive mechanisms, and we then posit that these mechanisms correspond to the general ability to solve certain classes of problems in a generalizable way. For instance, we suggest that it is only through inference control schemata which adapt based on experience that uncertain inference-based AI systems can learn to consistently solve Piagetan concrete-operational tasks in a way that provides knowledge suitable for further generalization. However, it may be that minds using hard-wired inference control schemata (typical of the infantile stage) can still solve some Piagetan concrete-operational tasks, though most solutions to such tasks obtained in this way will be “brittle” and not easily generalizable to other tasks using infantile cognition.

VI. INFANTILE COGNITION

One example of a cognitive task at the borderline between infantile and concrete cognition is learning object permanence, a problem discussed in a Novamente/AGI-SIM context in [46]. Another example is the learning of word-object associations: e.g. learning that when the word “ball” is uttered in various contexts (“Get me the ball,” “That’s a nice ball,” etc.) it generally refers to a certain type of object.

The key point regarding these “infantile” inference problems, from the Novamente perspective, is that assuming one provides the inference system with an appropriate set of perceptual and motor ConceptNodes and SchemaNodes, the chains of inference involved are short. They involve about a dozen inferences, and this means that the search tree of possible PLN inference rules walked by the PLN backward-chainer is relatively shallow. Sophisticated inference control is not required: standard AI heuristics are sufficient.

In short, textbook narrow-AI reasoning methods, utilized with appropriate uncertainty-savvy truth value formulas and coupled with appropriate representations of perceptual and motor inputs and outputs, correspond roughly to Piaget’s infantile stage of cognition. The simplistic approach of these narrow-AI methods may be viewed as a method of creating building blocks for subsequent, more sophisticated heuristics.

VII. CONSERVATION OF NUMBER

Next, as an example of a learning problem classically categorized within Piaget’s concrete-operational phase, we consider a “conservation laws” problem, discussed in [1] in the context of software that solves the problem using (logic-based and neural net) narrow-AI techniques.



Fig. 1: Conservation of Number

Conservation is the idea that a quantity remains the same despite changes in appearance. If you show a child some objects (Fig. 1) and then spread them out, an infantile mind will focus on the spread, and believe that there are now more objects than before, whereas a concrete-operational mind will understand that the quantity of objects has not changed.

Conservation of number seems very simple, but from a developmental perspective it is actually rather difficult. “Solutions” like those given in [1] that use neural networks or customized logical rule-bases to find specialized solutions that solve only this problem fail to fully address the issue, because these solutions don’t create knowledge adequate to aid with the solution of related sorts of problems.

We hypothesize that this problem is hard enough that for an inference-based AI system to solve it in a developmentally useful way, its inferences must be guided by meta-inferential lessons learned from prior similar

problems. When approaching a number conservation problem, for example, a reasoning system might draw upon past experience with set-size problems (which may be trial-and-error experience). This is not a simple “machine learning” approach whose scope is restricted to the current problem, but rather a heuristically guided approach which (a) aggregates information from prior experience to guide solution formulation for the problem at hand, and (b) adds the present experience to the set of relevant information about quantification problems for future refinement of thinking.

For instance, a very simple context-specific heuristic that a system might learn would be: “When evaluating the truth value of a statement related to the number of objects in a set, it is generally not that useful to explore branches of the backwards-chaining search tree that contain relationships regarding the sizes, masses, or other physical properties of the objects in the set.” This heuristic itself may go a long way toward guiding an inference process toward a correct solution to the problem—but it is not something that a mind needs to know “a priori.” A concrete-operational stage mind may learn this by data-mining prior instances of inferences involving sizes of sets. Without such experience-based heuristics, the search tree for such a problem will likely be unacceptably large. Even if it is “solvable” without such heuristics, the solutions found may be overly fit to the particular problem and not usefully generalizable.

VIII. THEORY OF MIND

Another learning problem that is typically classed in the Piagetan concrete-operational stage is “theory of mind” – which means, in this context, fully understanding the fact that others have memories, perceptions and experiences.

Consider this experiment: a preoperational child is shown her favorite “Dora the Explorer” DVD box. Asked what show she’s about to see, she’ll answer “Dora.” However, when her parent plays the disc, it’s “Spongebob Squarepants.” If you then ask her what show her friend will expect when given the “Dora” DVD box, she will respond “Spongebob” although she just answered “Dora” for herself. A child lacking a theory of mind can not reason through what someone else would think given knowledge other than her own current knowledge. Knowledge of self is intrinsically related to the ability to differentiate oneself from others, and this ability may not be fully developed at birth.

Several theorists [47,48], based in part on experimental work with autistic children, perceive theory of mind as embodied in an innate module of the mind activated at a certain developmental stage (or not, if damaged). While we consider this possible, we caution against adopting a simplistic view of the “innate vs. acquired” dichotomy: if there is innateness it may take the form of an innate predisposition to certain sorts of learning [49].

Davidson [50], Dennett [51] and others support the common belief that theory of mind is dependent upon linguistic ability. A major challenge to this prevailing

philosophical stance came from Premack and Woodruff [49] who postulated that prelinguistic primates do indeed exhibit “theory of mind” behavior. While Premack and Woodruff’s experiment itself has been challenged [52], their general result has been bolstered by follow-up work showing similar results such as [53]. It seems to us that while theory of mind depends on many of the same inferential capabilities as language learning, it is not intrinsically dependent on the latter.

There is a school of thought often called the *Theory Theory* [54]-[55]-[56] holding that a child’s understanding of mind is best understood in terms of the process of iteratively formulating and refuting a series of naïve theories about others. Alternately, Gordon [57] postulates that theory of mind is related to the ability to run cognitive simulations of others’ minds using one’s own mind as a model. We suggest that these two approaches are actually quite harmonious with one another. In an uncertain AI context, both theories and simulations are grounded in collections of uncertain implications, which may be assembled in context-appropriate ways to form theoretical conclusions or to drive simulations. Even if there is a special “mind-simulator” dynamic in the human brain that carries out simulations of other minds in a manner fundamentally different from explicit inferential theorizing, the inputs to and the behavior of this simulator may take inferential form, so that the simulator is in essence a way of efficiently and implicitly producing uncertain inferential conclusions from uncertain premises.

The details via which a Novamente system should be able to develop theory of mind in the AGI-SIM world have been articulated in detail, though practical learning experiments in this direction have not yet been done. We have not yet explored the possibility of giving Novamente a special “mind-simulator” component, though this would be possible; instead we have initially been pursuing a more purely inferential approach.

First, it is very simple for a Novamente system to learn patterns such as “If I rotated by pi radians, I would see the yellow block.” And it’s not a big leap for PLN to go from this to the recognition that “You look like me, and you’re rotated by pi radians relative to my orientation, therefore you probably see the yellow block.” The only nontrivial aspect here is the “you look like me” premise.

Recognizing “embodied agent” as a category, however, is a problem fairly similar to recognizing “block” or “insect” or “daisy” as a category. Since the Novamente agent can perceive most parts of its own “robot” body—its arms, its legs, etc.—it should be easy for the agent to figure out that physical objects like these look different depending upon its distance from them and its angle of observation. From this it should not be that difficult for the agent to understand that it is naturally grouped together with other embodied agents (like its teacher), not with blocks or bugs.

The only other major ingredient needed to enable theory of mind is “reflection”—the ability of the system to explicitly recognize the existence of knowledge in its own mind (note that this term “reflection” is not the same as our

proposed “reflexive” stage of cognitive development). This exists automatically in Novamente, via the built-in vocabulary of elementary procedures supplied for use within SchemaNodes (specifically, the *atTime* and *TruthValue* operators). Observing that “at time T, the weight of evidence of the link L increased from zero” is basically equivalent to observing that the link L was created at time T.

Then, the system may reason, for example, as follows (using a combination of several PLN rules including the above-given deduction rule):

Implication

My eye is facing a block and it is not dark

A relationship is created describing the block’s color

Similarity

My body

My teacher’s body

|-

Implication

My teacher’s eye is facing a block and it is not dark

A relationship is created describing the block’s color

This sort of inference is the essence of Piagetan “theory of mind.” Note that in both of these implications the created relationship is represented as a variable rather than a specific relationship. The cognitive leap is that in the latter case the relationship actually exists in the teacher’s implicitly hypothesized mind, rather than in Novamente’s mind. No explicit hypothesis or model of the teacher’s mind need be created in order to form this implication—the hypothesis is created implicitly via inferential abstraction. Yet, a collection of implications of this nature may be used via an uncertain reasoning system like PLN to create theories and simulations suitable to guide complex inferences about other minds.

From the perspective of developmental stages, the key point here is that in a Novamente context this sort of inference is too complex to be viably carried out via simple inference heuristics. This particular example must be done via forward chaining, since the big leap is to actually think of forming the implication that concludes inference. But there are simply too many combinations of relationships involving Novamente’s eye, body, and so forth for the PLN component to viably explore all of them via standard forward-chaining heuristics. Experience-guided heuristics are needed, such as the heuristic that if physical objects A and B are generally physically and functionally similar, and there is a relationship involving some part of A and some physical object R, it may be useful to look for similar relationships involving an analogous part of B and objects similar to R. This kind of heuristic may be learned by experience—and the masterful deployment of such heuristics to guide inference is what we hypothesize to characterize the concrete stage of development. The “concreteness” comes from the fact that inference control is guided by analogies to prior similar situations.

IX. SYSTEMATIC EXPERIMENTATION

The Piagetan formal phase is a particularly subtle one from the perspective of uncertain inference. In a sense, AI inference engines already have strong capability for formal reasoning built in. Ironically, however, no existing inference engine is capable of deploying its reasoning rules in a powerfully effective way, and this is because of the lack of inference control heuristics adequate for controlling abstract formal reasoning. These heuristics are what arise during Piaget's formal stage, and we propose that in the content of uncertain inference systems, they involve the application of inference itself to the problem of refining inference control.

A problem commonly used to illustrate the difference between the Piagetan concrete operational and formal stages is that of figuring out the rules for making pendulums swing quickly versus slowly [10]. If you ask a child in the formal stage to solve this problem, she may proceed to do a number of experiments, e.g. build a long string with a light weight, a long string with a heavy weight, a short string with a light weight and a short string with a heavy weight. Through these experiments she may determine that a short string leads to a fast swing, a long string leads to a slow swing, and the weight doesn't matter at all.

The role of experiments like this, which test "extreme cases," is to make cognition easier. The formal-stage mind tries to map a concrete situation onto a maximally simple and manipulable set of abstract propositions, and then reason based on these. Doing this, however, requires an automated and instinctive understanding of the reasoning process itself. The above-described experiments are good ones for solving the pendulum problem because they provide data that is very easy to reason about. From the perspective of uncertain inference systems, this is the key characteristic of the formal stage: formal cognition approaches problems in a way explicitly calculated to yield tractable inferences.

Note that this is quite different from saying that formal cognition involves abstractions and advanced logic. In an uncertain logic-based AI system, even infantile cognition may involve these--the difference lies in the level of inference control, which in the infantile stage is simplistic and hard-wired, but in the formal stage is based on an understanding of what sorts of inputs lead to tractable inference in a given context.

X. CORRECTION OF INFERENCE BIASES

Finally, we will briefly allude to an example of what we've called the "reflexive" stage in inference. Recall that this is a stage beyond Piaget's formal stage, reflecting the concerns of [25,28-31] that the Piagetan hierarchy ignores the ongoing development of cognition into adulthood.

Highly intelligent and self-aware adults may carry out reflexive cognition by explicitly reflecting upon their own inference processes and trying to improve them. An

example is the intelligent improvement of uncertain-truth-value-manipulation formulas. It is well demonstrated that even educated humans typically make numerous errors in probabilistic reasoning [57,58]. Most people don't realize it and continue to systematically make these errors throughout their lives. However, a small percentage of individuals make an explicit effort to increase their accuracy in making probabilistic judgments by consciously endeavoring to internalize the rules of probabilistic inference into their automated cognition processes.

The same sort of issue exists even in an AI system such as Novamente which is explicitly based on probabilistic reasoning. PLN is founded on probability theory, but also contains a variety of heuristic assumptions that inevitably introduce a certain amount of error into its inferences. For example, the probabilistic deduction formula mentioned above embodies a heuristic independence assumption. Thus PLN contains an alternate deduction formula called the "concept geometry formula" [41] that is better in some contexts, based on the assumption that ConceptNodes embody concepts that are roughly spherically-shaped in attribute space. A highly advanced Novamente system could potentially augment the independence-based and concept-geometry-based deduction formulas with additional formulas of its own derivation, optimized to minimize error in various contexts. This is a simple and straightforward example of reflexive cognition--it illustrates the power accessible to a cognitive system that has formalized and reflected upon its own inference processes, and that possesses at least some capability to modify these.

XI. CONCLUSION

AI systems must *learn*, but they must also *develop*: and development in this sense takes place over a longer time scale than learning, and involves more fundamental changes in cognitive operation. Understanding the development of cognition is equally as important to AI as understanding the nature of cognition at any particular stage.

We have proposed a novel approach to defining developmental stages, in which internal properties of inference control systems are correlated with external learning capabilities, and have fleshed out the approach via giving a series of specific examples related to the Novamente AI Engine and the AGI-SIM world. Our future work with Novamente will involve teaching it to perform behaviors in the AGI-SIM world, progressing gradually through the developmental stages described here, using examples such as those given. Finally, we suspect that this approach to developmental psychology also has relevance beyond Novamente--most directly to other uncertain inference-based AI systems, and perhaps to developmental psychology in general.

REFERENCES

- [1] T. Shultz. *Computational Developmental Psychology*. Cambridge, MA: MIT Press, 2003.

- [2] E. Thelen and L. Smith. *A Dynamic Systems Approach to the Development of Cognition and Action*. Cambridge, MA: MIT Press, 1994.
- [3] J. A.S. Kelso. *Dynamic Patterns: The Self-Organization of Brain and Behavior*. Cambridge, MA: MIT Press, 1995.
- [4] M. Howe and F. Rabinowitz. "Dynamic modeling, chaos, and cognitive development." *Journal of Experimental Child Psychology*, 1994, 58, 184-199.
- [5] B. Goertzel, C. Pennachin, A. Senna, T. Maia, G. Lamacie. "Novamente: an integrative architecture for Artificial General Intelligence." *Proceedings of IJCAI 2003 Workshop on Cognitive Modeling of Agents*. Acapulco, Mexico, 2004.
- [6] M. Looks, B. Goertzel and C. Pennachin. "Novamente: an integrative architecture for Artificial General Intelligence." *Proceedings of AAAI 2004 Symposium on Achieving Human-Level AI via Integrated Systems and Research*, Washington DC, 2004.
- [7] B. Goertzel and Cassio Pennachin. "The Novamente AI engine." *Artificial General Intelligence*, Ed. by B. Goertzel and C. Pennachin, New York: Springer Verlag, 2005.
- [8] B. Goertzel. *The Foundations of Emergent Digital Intelligence*. In preparation.
- [9] J. Santore and S. Shapiro. "Crystal Cassie: Use of a 3-D gaming environment for a Cognitive Agent." In R. Sun, Ed., *Papers of the IJCAI 2003 Workshop on Cognitive Modeling of Agents and Multi-Agent Interactions*, IJCAI, Acapulco, Mexico, Aug. 9, 2003, 84-91.
- [10] B. Inhelder and J. Piaget. *The Growth of Logical Thinking from Childhood to Adolescence*. New York: Basic Books, 1958.
- [11] J. Piaget. *Structuralism*. New York: Harper & Row, 1970.
- [12] J. Piaget. *Biology and Knowledge*. Chicago: University of Chicago Press, 1971.
- [13] J. Piaget. "Piaget's theory." In P. Mussen (ed). *Handbook of Child Psychology*. 4th edition. Vol. 1. New York: Wiley, 1983
- [14] J. Piaget. *Studies in Reflecting Abstraction*. Hove, UK: Psychology Press, 2001.
- [15] S. Oppen and H. Ginsburg. *Piaget's Theory of Development*. New York: Prentice-Hall, 1987.
- [16] L. Vygotsky. *Mind in Society: The Development of Higher Psychological Processes*. Cambridge, MA: Harvard University Press, 1978.
- [17] L. Vygotsky. *Thought and Language*. Cambridge, MA: MIT Press, 1986.
- [18] R., Gagne, L. Briggs, and W. Walter. *Principles of Instructional Design*, Fort Worth: Harcourt Brace Jovanovich College Publishers, 1992.
- [19] J. Gibbs, "Kohlberg's moral stage theory: a Piagetian revision." *Human Development*, 1978, 22, 89-112.
- [20] J. C. Gibbs, K. F. Widaman, and A. Colby. "Sociomoral reflection: Construction and validation of group-administrable measures." *Child Development*, 1982, 53, 895-910.
- [21] J. Broughton. "Not beyond formal operations, but beyond Piaget." In M. Commons, F. Richards, and C. Armon (Eds.), *Beyond Formal Operations: Late Adolescent and Adult Cognitive Development*. New York: Praeger, 1984, 395-411.
- [22] M.L. Commons and A. Pekker, *Hierarchical Complexity: A Formal Theory*, unpublished manuscript, submitted to J. Math. Psych., 2005.
- [23] J. Pascual-Leone and J. Smith. "The encoding and decoding of symbols by children: a new experimental paradigm and neo-Piagetan model." *Journal of Experimental Child Psychology*, 1969, 8, 328-355.
- [24] R. Gelman, E. Meck and S. Merkin. "Young children's numerical competence." *Cognitive Development*, 1986, 1, 1-29.
- [25] K. Fischer. "A theory of cognitive development: control and construction of hierarchies of skills." *Psychological Review*, 1980, 87, 477-531.
- [26] R. Case. *Intellectual development: Birth to adulthood*. New York: Academic Press, 1985.
- [27] M. Bickhard. "Piaget on variation and selection models: Structuralism, logical necessity, and interactivism." *Human Development*, 1988, 31, 274-312.
- [28] Arlin, P.K. "Cognitive development in adulthood: A fifth stage?" *Developmental Psychology*, 1975, 11, 602-606.
- [29] M. Commons, F. Richards and D. Kuhn. "Systematic and metasystematic reasoning: a case for a level of reasoning beyond Piaget's formal operations." *Child Development*, 1982, 53, 1058-1069.
- [30] K. Riegel. "Dialectic operations: the final phase of cognitive development." *Human Development*, 1973, 16, 346-370.
- [31] H. Marchand. "Reflections on PostFormal Thought." In *The Genetic Epistemologist*, vol. 29, no. 3, 2001.
- [32] K. Kitchener and P. King. "Reflective judgement: ten years of research." In M. Commons, et. Al. (Eds.), *Beyond Formal Operations Vol. 2: Models and Methods in the Study of Adolescent and Adult Thought*. New York: Praeger, 1990, 63-78.
- [33] Commons, M., Trudeau, E.J., Stein, S.A., Richards, F.A., & Krause, S.R. "Hierarchical complexity of tasks shows the existence of developmental stages." *Developmental Review*, 1998, 18, 237-278.
- [34] G. Engelbretsen and F. Sommers. "An invitation to formal reasoning." *The Logic of Terms*. Aldershot: Ashgate, 2000.
- [35] L. Zadeh. Fuzzy sets as a basis for a theory of possibility. *Fuzzy Sets and System*, 1978, 1:3-28.
- [36] I. Good. *The Estimation of Probabilities*. Cambridge, MA: MIT Press, 1986.
- [37] R. Fung and C. Chong. Metaprobability and Dempster-Shafer in evidential reasoning. In L. Kanal and J. Lemmer (Eds.), *Uncertainty in Artificial Intelligence*, Amsterdam, Holland, 1986, 295-302.
- [38] G. Shafer. *A Mathematical Theory of Evidence*. Princeton, NJ: Princeton University Press, 1976.
- [39] H. Kyburg. Bayesian and non-Bayesian evidential updating. *Artificial Intelligence*, 1987, 31:271-293.
- [40] P. Wang. *Non-Axiomatic Reasoning System*. PhD Thesis, Bloomington, IN: Indiana University, 1995.
- [41] B. Goertzel, I. F. Goertzel, M. Iklé, A. Heljakka. *Probabilistic Logic Networks* In preparation.
- [42] D. MacKenzie. The automation of proof: A historical and sociological exploration. *IEEE Annals of the History of Computing*, 1995, 17(3), 7-29.
- [43] B. Bollobas. *Modern Graph Theory*. New York: Springer 1998
- [44] P. Walley. *Statistical Reasoning with Imprecise Probabilities*. New York: Chapman and Hall, 1991.
- [45] S. Russell and P. Norvig. *Artificial Intelligence: A Modern Approach*, Englewood Cliffs, NJ: Prentice Hall, 1995.
- [46] B. Goertzel. Patterns, Hypergraphs and Embodied General Intelligence. Submitted to *Workshop on Human-Level Intelligence, WCCI 2006*, Vancouver.
- [47] S. Baron-Cohen. *Mindblindness: An Essay on Autism and Theory of Mind*. Cambridge, MA: MIT Press, 1995.
- [48] J. Fodor. *The Elm and the Expert*, Cambridge, MA: Bradford Books, 1994.
- [49] J. Elman, E. Bates, M. Johnson, A. Karmiloff-Smith, D. Parisi, K. Plunkett. *Rethinking Innateness: A Connectionist Perspective on Development*. MIT Press, 1997.
- [50] D. Davidson. *Inquiries into Truth and Interpretation*, Oxford: Oxford University Press, 1984.
- [51] D. Dennett. *The Intentional Stance*, Cambridge, MA: MIT Press, 1987.
- [52] D. Premack and G. Woodruff. "Does the chimpanzee have a theory of mind?" *Behavioral and Brain Sciences*, 1978 1, 515-526.
- [53] M. Tomasello and J. Call. *Primate Cognition*. New York: Oxford University Press, 1997.
- [54] R. W. Byrne and A. Whiten. *Machiavellian Intelligence*. Clarendon Press, 1988.
- [55] S. Carey. *Conceptual Change in Childhood*, Cambridge, MA: MIT Press, 1985.
- [56] H. Wellman. *The Child's Theory of Mind*, Cambridge, MA: MIT Press, 1990.
- [57] R. Gordon. Folk Psychology as Simulation. *Mind and Language*, 1986, 1, 158-171.
- [58] P. Wason. Reasoning. In B. M. Foss (Ed.), *New Horizons in Psychology*. Middlesex, UK: Penguin, 1966
- [59] T. Gilovich, D. Griffin & D. Kahneman, (Eds.). *Heuristics and biases: The psychology of intuitive judgment*. Cambridge, UK: Cambridge University Press, 2002