Perceptual Memory and Learning: Recognizing, Categorizing, and Relating

Stan Franklin

Computer Science Division &
Institute for Intelligent Systems
The University of Memphis
Memphis, TN 38152, USA
franklin@memphis.edu

Abstract

In this position paper we attempt to derive an architecture and mechanism for perceptual memory and learning for software agents and robots from what is known, or believed, about the same faculties in human and other animal cognition. Based on that of the IDA model of Global Workspace Theory, a conceptual and computational model of cognition, this architecture, together with its mechanisms, offers the real possibility of autonomous software agents and robots learning their own ontologies during a developmental period. Thus the onerous chore of designing and implementing such an ontology can be avoided.

Premises

In particular we want to base the design on the following premises:

1. *Perceptual memory, that is the ability to interpret incoming stimuli by recognizing individuals, by categorizing them, and by noting relationships between such individuals and categories, is ubiquitous among animal species, as is the learning of these facilities* (Bitterman 1965).

    Animals of all sorts can identify food sources, potential mates, potential predators, etc. Pigeons have been taught to categorize using such concepts as tree, fish, and human, some well outside of their evolutionary background (Herrnstein 1984). Honey bees have been taught to identify human letters independently of size, color, position or font (Gould & Gould 1988). An African Grey Parrot can identify such features as size, number, color, and material of objects or sets of objects that he has never been seen before (Pepperberg 1990).

2. *Perceptual memory is evolutionarily older than semantic memory in humans, and has its own, distinct mechanism* (Franklin et al. in review).

    There are developmental arguments for a distinct mechanism for perceptual memory; infants who have not yet developed object permanence (any episodic memory) are quite able to recognize and categorize (Mandler 2000). Other arguments come from studies of human amnesiacs with significant loss of new declarative memory, but mostly intact perceptual memory and learning (Gabrieli et al. 1990, Fahle & Daum 2002). Perhaps the most convincing argument comes from experiments with rats in a radial arm maze. With four arms baited and four not (with none restocked), normal rats learn to recognize which arms to search (perceptual memory) and remember which arms they have already fed in (episodic memory) so as not to search there a second time. Rats with their hippocampal systems excised lose their episodic memory but retain perceptual memory, again arguing for distinct mechanisms (Olton et al. 1979).

3. *Conscious awareness is sufficient for perceptual learning* (Baars 1988, Hobson & Stickgold 1994).

    This premise may be the most controversial of the lot, with people arguing that implicit learning, say for example, learning to distinguish well-formed strings from a finite grammar without learning the rules, doesn't require consciousness. Subjects learn some implicit version of the grammar by being consciously exposed to well- and ill-formed strings and being told which is which. Consciousness is required for this exposure.

4. *Perceptual learning is facilitated by feelings and emotions* (Franklin & McCaulley 2004).

    In particular, the learning rate varies with arousal (Yerkes & Dodson 1908).

5. *Perceptual learning occurs easily and rapidly, but decays according to an inverse sigmoid function; new, weak memories decay extraordinarily rapidly, while saturated perceptual memories may last for many decades* (Franklin et al. in review).

6. Preconscious perception is the first step in a continually cascading series of cognitive cycles by means of which every animal samples its environment and acts on it (Baars & Franklin 2003, Franklin et al. in review).

With these premises in hand, we will outline the design of a perceptual memory and learning mechanism within the conceptual and computational IDA model of cognition, which is derived from the IDA software agent (Franklin & Graesser 1997).

The IDA Model

IDA provides a conceptual (and computational) model of cognition (Franklin 2000, 2001b) partially implemented as a software agent (Franklin & Graesser 1997). The implemented IDA "lives" on a computer system with connections to the Internet and various databases, and does personnel work for the US Navy, performing all the specific personnel tasks of a human (Franklin 2001a). In particular, IDA negotiates with sailors in natural language,

deliberates, and makes voluntary action selections in the process of finding new jobs for sailors at the end of their current tour of duty. IDA completely automates the work of Navy personnel agents (detailers).

The IDA model implements and fleshes out Global Workspace theory (Baars 1988, 2002), which suggests that conscious events involve widespread distribution of focal information needed to recruit neuronal resources for problem solving. The IDA implementation of GW theory yields a fine-grained functional account of the steps involved in perception, several kinds of memory, consciousness, context setting, and action selection. Cognitive processing in IDA consists of continually repeated traversals through the steps of a *cognitive cycle* (Baars & Franklin 2003, Franklin et al. in review), as described below.

The IDA architecture includes modules for perception (Zhang, et al. 1998), various types of memory (Anwar and Franklin. 2003, Franklin et al. in review), "consciousness" (Bogner, Ramamurthy and Franklin. 2000), action selection (Negatu and Franklin. 2002), constraint satisfaction (Kelemen, Liang, and Franklin. 2002), deliberation (Franklin 2000a), and volition (Franklin 2000a). The mechanisms of these modules are derived from several different "new AI" sources (Hofstadter and Mitchell. 1994, Jackson 1987, Maes 1989).

IDA senses strings of characters from email messages and databases, and negotiates with sailors via email. The computational IDA is a running software agent that has been tested and demonstrated to the satisfaction of the Navy. Detailers observing the testing often commented that "IDA thinks like I do."

The running, computational IDA software agent is almost entirely handcrafted. There is essentially no learning. However, in addition to the computational model, we will also speak of the conceptual IDA model, which includes additional capabilities that have been designed but not implemented, including mechanisms for feelings and emotions, and the mechanism for perceptual learning, the primary focus of this paper.

The IDA conceptual model contains several different memory systems. Perceptual memory enables identification, recognition and categorization, including of feelings, as well as relationships. Working memory provides preconscious buffers as a workspace for internal activities. Transient episodic memory is a content-addressable associative memory with a moderately fast decay rate. It is to be distinguished from autobiographical memory, a part of declarative memory. Procedural memory is long-term memory for skills. The material that follows describes perceptual learning as envisioned in the IDA conceptual model.

The IDA Technology

The IDA Technology is based on a number of highly connected modules each built on its distinct mechanism. Most of these are up and running. A few are still being developed, and a couple are designed but not yet implemented. Following Hofstadter's terminology (see below) a codelet is a special purpose, relatively independent, mini-agent typically implemented as a small piece of code running as a separate thread. IDA depends heavily on such codelets for almost every module. In what follows we will encounter several different types of codelets such as perceptual codelets, attention codelets, information codelets, behavior codelets and expectation codelets. Many codelets play the role of demons (as in an operating system) waiting patiently for the conditions under which they can act. Some codelets subserve some higher-level construct, while others act completely independently. Neurally, they can be thought of as cell assemblies or neuronal groups (Edelman1987, Edelman and Tononi 2000).

In this section we describe several of the IDA modules that would play a role in perceptual memory and learning.

Perception

IDA senses only strings of characters. Perception consists mostly of processing incoming email messages in natural language (Zulandt Schneider et al. 2001). In sufficiently narrow domains, natural language understanding may be achieved via an analysis of surface features without the use of a traditional symbolic parser (Jurafsky & Martin 2000). Allen describes this approach as complex, template-based matching (1995).

IDA's perceptual module has been implemented (as yet without learning) as a Copycat-like architecture with perceptual codelets that are triggered by surface features and a slipnet (Hofstadter & Mitchell 1994), a semantic net that passes activation. The slipnet stores domain knowledge. In addition there's a pool of perceptual codelets specialized for recognizing particular pieces of text, and production templates used by codelets for building and verifying understanding. Together they constitute an integrated perceptual system for IDA, allowing her to recognize, categorize and understand.

An example will illustrate what is claimed by the word "understand" as used in the previous sentence. A clerical worker sending out an email announcement of an upcoming seminar on Compact Operators on Banach Spaces can be said to have understood the organizer's request that this be done even though he or she has no idea of what a Banach space is much less what compact operators on them are. In most cases it would likely require several person years of diligent effort to impart such knowledge. Nonetheless, the clerical worker understands the request at a level sufficient for him or her to get out the announcement. In the same way IDA understands incoming email messages well enough to do all the things needed to be done with them. An expanded form of this argument can be found in *Artificial Minds* (Franklin 1995). Glenberg also makes a similar argument (1997).

An underlying assumption motivates our design decisions about perception. Suppose, for example, that IDA receives a message from a sailor saying that his projected rotation date is approaching and asking that a job

be found for him. The perception module would recognize the sailor's name and social security number, and that the message is of the please-find-job type. This information would then be written to the workspace. The general principle here is that the contents of perception are written to working memory before becoming conscious.

Workspace

IDA solves routine problems with novel content. This novel content goes into her workspace, which roughly plays the same role as the preconscious buffers of human working memory. Perceptual codelets write to the workspace as do other, more internal codelets. Quite a number of codelets, including attention codelets (see below) watch what's written in the workspace in order to react to it. Part, but not all, the workspace, called the *focus[1]*, by Kanerva (1988) is set aside as an interface with episodic memory (EM). Retrievals from EM are made with cues taken from the focus and the resulting associations are written to other registers in the focus. The contents of still other registers in the focus are stored in (written to) EM. Items in the workspace decay over time, and may be overwritten. Not all of the contents of the workspace eventually make their way into consciousness.

Episodic memory

IDA employs sparse distributed memory (SDM) as her major episodic memory (Kanerva 1988, Anwar & Franklin 2003). SDM is a content addressable memory that, in many ways, is an ideal computational mechanism for use as an EM. Any item written to the workspace cues a retrieval from EM, returning prior activity associated with the current entry. EM is accessed as soon as information reaches the workspace, and the retrieved associations will be also written to the workspace.

At a given moment IDA's workspace may contain, ready for use, a current entry from perception or elsewhere, prior entries in various states of decay, and associations instigated by the current entry, i.e. activated elements of EM. IDA's workspace thus consists of both short-term working memory and something like the long-term working memory of Ericsson and Kintsch (1995).

"Consciousness" mechanism

The apparatus for "consciousness"[2] consists of a coalition manager, a spotlight controller, a broadcast manager, and a collection of attention codelets whose job it is to bring appropriate contents to "consciousness" (Bogner et al. 2000). Each attention codelet keeps a watchful eye out for some particular occurrence that might call for "conscious" intervention. In most cases an attention codelet is watching

the workspace, which will likely contain both perceptual information and data created internally, the products of "thoughts." Upon encountering such a situation, the appropriate attention codelet will form a coalition with the small number of information codelets that carry the information describing the situation. This association should lead to the collection of this small number of information codelets, together with the attention codelet that collected them, becoming a coalition. Codelets also have activations. The attention codelet increases its activation in order that the coalition, if one is formed, might compete for the spotlight of "consciousness". Upon winning the competition, the contents of the coalition is then broadcast to all codelets. If or when successful, its contents will be broadcast.

Action selection (decision making)

IDA depends on a behavior net (Maes 1989, Negatu & Franklin 2002) for high-level action selection in the service of built-in drives. Several distinct drives, operating in parallel, vary in urgency as time passes and the environment changes. Behaviors are typically mid-level actions, many depending on several behavior codelets for their execution. A behavior net is composed of behaviors, corresponding to goal contexts in GW theory, and their various links. A behavior looks very much like a production rule, having preconditions as well as additions and deletions, but at a higher level of abstraction, often requiring the efforts of several behavior codelets to effect its action. Each behavior occupies a node in a digraph.

As in connectionist models (McClelland et al. 1986), this digraph spreads activation. The activation comes from that stored in the behaviors themselves, from the environment, from drives, and from internal states. The more relevant a behavior is to the current situation, the more activation it receives from the environment. Each drive awards activation to those behaviors that will satisfy it. Certain internal states of the agent can also send activation to the behavior net. One example might be activation from a coalition of codelets responding to a "conscious" broadcast. Activation spreads from behavior to behavior along both excitatory and inhibitory links and a behavior is chosen to execute based on activation. The behavior net produces flexible, tunable action selection for IDA. As is widely recognized in humans the hierarchy of goal contexts is fueled at the top by drives, that is, by primitive motivators implemented as feelings and emotions, and at the bottom by input from the environment, both external and internal.

Each "conscious" broadcast is received by appropriate behavior codelets who know to instantiate a behavior stream in the behavior net for dealing with the current situation. They also bind appropriate variables, and send activation to appropriate behaviors. If or when a particular behavior is chosen to be executed, behavior codelets associated with it each perform its assigned task.

---

[1] Not to be confused with focus as in focus of attention, an entirely different concept.
[2] The scare quotes signal that no claim of subjective consciousness is being made, only for functional consciousness (Franklin 2003).

There's much more to the IDA architecture and mechanisms, but this is all that space will allow.

The Cognitive Cycle

The IDA model suggests a number of more specialized roles for feelings in cognition, all combining to produce motivations and to facilitate learning. Here we describe the conceptual IDA's cognitive cycle, most, but not all, of which has been implemented.

1. *Perception. Sensory stimuli, external or internal, are received and interpreted by perception producing meaning.* Note that this stage is unconscious.
   a. Early perception: *Input arrives through senses. Specialized perception codelets descend on the input. Those that find features relevant to their specialty activate appropriate nodes in the slipnet (a semantic net with activation).*
   This perceptual memory system identifies pertinent feeling/emotions are along with objects, categories and their relations.
2. *Percept to Preconscious Buffer. The percept, including some of the data plus the meaning, is stored in preconscious buffers of IDA's working memory.*
   In humans, these buffers may involve visuo-spatial, phonological, and other kinds of information. Feelings/emotions are part of the preconscious percept.
3. *Local Associations. Using the incoming percept and the residual contents of the preconscious buffers, including emotional content, as cues, local associations are automatically retrieved from transient episodic memory (TEM) and from declarative memory.*
   The contents of the preconscious buffers together with the retrieved local associations from TEM and declarative memory, roughly correspond to Ericsson and Kintsch's long-term working memory (1995) and to Baddeley's episodic buffer (2000). These local associations include records of the agent's past feelings/emotions, and actions, in associated situations.
4. *Competition for Consciousness. Attention codelets view long-term working memory, and bring relevant, urgent, or insistent events to consciousness,. Some of them gather information, form coalitions and actively compete for access to consciousness.*
   The competition may also include such coalitions from a recently previous cycle. Present and past feelings/emotions influence this competition for consciousness. Strong affective content strengthens a coalition's chances of coming to consciousness.
5. *Conscious Broadcast. A coalition of codelets, typically an attention codelet and its covey of related information codelets carrying content, gains access to the global workspace and has its contents broadcast.*
   In humans, this broadcast is hypothesized to correspond to phenomenal consciousness. *The conscious broadcast contains the entire content of consciousness including the affective portions.* The contents of perceptual memory are

updated in light of the current contents of consciousness, *including feelings/emotions*, as well as objects, categories and relations. *The stronger the affect, the stronger the encoding in memory.* Transient episodic memory is also updated with the current contents of consciousness, *including feelings/emotions*, as events. *The stronger the affect, the stronger the encoding in memory.* (At recurring times not part of a cognitive cycle, the contents of transient episodic memory are consolidated into long-term declarative memory.) Procedural memory (recent actions) is updated (reinforced) with the strength of the reinforcement influenced by the strength of the affect.

6. *Recruitment of Resources. Relevant behavior codelets respond to the conscious broadcast. These are typically codelets whose variables can be bound from information in the conscious broadcast.*
   If the successful attention codelet was an expectation codelet calling attention to an unexpected result from a previous action, the responding codelets may be those that can help to rectify the unexpected situation. Thus consciousness solves the relevancy problem in recruiting resources. The affective content (feelings/emotions) together with the cognitive content, help to attract relevant resources (processors, neural assemblies) with which to deal with the current situation.
7. *Setting Goal Context Hierarchy. The recruited processors use the contents of consciousness, including feelings/emotions, to instantiate new goal context hierarchies, bind their variables, and increase their activation.*
   It is here that feelings and emotions most directly implement motivations by helping to instantiate and activate goal contexts, and by determining which terminal goal contexts receive activation. Other, environmental, conditions determine which of the earlier goal contexts receive additional activation.
8. *Action Chosen. The behavior net chooses a single behavior (goal context), perhaps from a just instantiated behavior stream or possibly from a previously active stream.*
   This selection is heavily influenced by activation passed to various behaviors influenced by the various feelings/emotions. The choice is also affected by the current situation, external and internal conditions, by the relationship between the behaviors, and by the residual activation values of various behaviors.
9. *Action Taken. The execution of a behavior (goal context) results in the behavior codelets performing their specialized tasks, which may have external or internal consequences, or both.*
   This is IDA taking an action. The acting codelets also include at least one expectation codelet (see Step 6) whose task it is to monitor the action and to try and bring to consciousness any failure in the expected results.
   We suspect that cognitive cycles occur five to ten times a second in humans, cascading so that some of the steps in adjacent cycles occur in parallel (Baars & Franklin 2003). Seriality is preserved in the conscious broadcasts.

Perceptual Memory and Learning

In this section we will describe in detail the structure and operation of the perceptual memory mechanism in the conceptual IDA model, as well as the procedures by which the memory is modified through learning.

As mentioned above, perceptual memory (PM) is implemented in the IDA model as a semantic net with passing activation, a slipnet. The nodes in the slipnet may represent feature detectors (perceptual codelets), individuals (e.g. a person or particular thing), a category (e.g. chair, woman, animal), a concept (e.g. democracy, justice), an idea (e.g. "please find me a job"), etc. Links in the slipnet represent relations between nodes, including category membership, category inclusion, and spatial, temporal or causal relations. Links can be excitatory or inhibitory. It's best to think of the slipnet as being feedforward in terms of conceptual depth (Hofstadter & Mitchell 1994), that is, the length of the shortest path from the periphery. This isn't completely correct, since lateral paths can also exist. Moving inward from the periphery means a link from the more specific to the more abstract, e.g. from the `Stan node` to the `man node` to the `human node` to the `animal node`.

Every autonomous agent senses its environment (Franklin & Graesser 1997), using such sensory modalities as vision, olfaction, audition, echolocation or, as in the software agent IDA, strings of characters. Thus each such agent, be it a software agent like IDA, or a human, must come equipped with primitive feature detectors, each with its receptive field among the appropriate sensory receptors. An example in humans might be a visual feature detector for an edge at a particular angle. These primitive feature detectors are, thus, directly connected to the agent's incoming senses, and in the IDA conceptual model constitute the nodes of minimal conceptual depth in the slipnet, its very periphery. Primitive feature detectors can combine to form more complex feature detectors, e.g. one to recognize a 'T' form. Feature detectors send activation along `is a feature of` links to item, category, concept and idea nodes further in the slipnet. As described in Step 1b of the cognitive cycle, the slipnet passes activation until it stabilizes. At this point those nodes and links with activation above threshold become part of IDA's percept, which gets passed along to the preconscious buffers of working memory in Step 2.

Some categories are essentially disjoint, e.g. `men` and `women`. Inhibitory `not` links run both ways between them. This results in a 'winner take all' stabilization with at most one of the nodes appearing in the percept connected to a particular person. For example, feature detectors for long hair and makeup might, in our culture, give the `woman` node an initial advantage, leading to its winning.

Each node in the slipnet (PM) has both a base-level activation and a current activation. Each link has only a base-level activation, some function of which acts as a weight on the link. Base-level activation is used for perceptual learning, while current activation results from the current exogenous or endogenous environment. Thus, current activation begins with primitive feature detectors and propagates inward to nodes at greater conceptual depth. Base-level and current activation are combined to arrive at the total activation of each node, upon which the percept is based.

Current activation decays quickly, so that it disappears completely within a small number of cognitive cycles, say in a second or two. Base-level activation decays in an inverse sigmoidal fashion, so that nodes or links with low base-level activations decay quite rapidly, while those with high, saturated base-level activations decay quite slowly, persisting in humans for decades. For example, I see the word "skunk" exceedingly rarely these days in my reading, but nonetheless, recognize it readily since it's base-level activation has saturated over the years.

Following Premise 3 above, perceptual learning in the IDA model occurs with consciousness. This learning is of two forms, the strengthening or weakening of the base-level activation of existing nodes and links, as well as the creation of new nodes and links. Any existing concept or relation that appears in the conscious broadcast (Step 5 of the cognitive cycle) has the base-level activation of its corresponding node or link strengthened as a function of the arousal of the agent (or weakened, depending on the valence) at the time of the broadcast.

A new individual item that comes to consciousness results in a new node being created, together with links into it from the feature detectors of its features. Such a new item gets to consciousness by means of some new item-attention codelet that notices a collection of active features in the percept without a common object of which they are features. Such a new item-attention codelet might be looking for such features as special contiguity, common motion, and persistence over time. If this attention codelet succeeds in bringing the resulting new item to consciousness, a node for it is created in PM by the perceptual learning mechanism.

New links (relations) occur similarly. Here's how a new category may be formed. If a similarity-attention codelet notices in long-term working memory (see Step 3 of the cognitive cycle) two items with several common features, and succeeds in bringing this similarity to consciousness, a new category is created by the perceptual learning mechanism with `is a` links into the new category from each of the items. Other new relations are learned into links in PM when other relation-noting-attention codelets succeed in bringing the new relations to consciousness, that is when the relation "pops into mind" or "occurs to me." The initial base-level activation of new nodes and links are assigned as a function of arousal at the time of the conscious broadcast (Step 5 of the cognitive cycle above).

One may object that all these new nodes and links, sometimes created as often as several times a second, might prove computationally intractable. But nature is often profligate; witness the vast numbers of acorns or sperm, so few of which come to any fruition. Here we have another example of such profligacy in perceptual learning. We are saved from computationally intractability by the

rapid decay of almost all of the new nodes and links, by virtue of their inverse sigmoidal decay curves. Only those new nodes and links that come to consciousness often and/or at high arousal levels have much chance of not decaying away. In the AI literature, a similar mechanism is referred to as generate and test. In the IDA conceptual model, perceptual learning generates trial nodes (combined feature detectors, individual items, categories, etc.) and links (relations), and rapidly discards those that don't quickly prove useful.

One can argue that the IDA conceptual model's slipnet every node is a category node. A primitive feature dectector can be though of as the category of the sensory receptors in its receptive field. A combined feature detector node can be thought of as the category of feature detectors of less conceptual depth of which it is composed, that is, with `is-a-feature-of` links into it. A individual item node can be thought of as the category of its feature detectors. And, nodes can represent categories of categories.

In keeping with Barsalou's Perceptual Symbol Systems (1999), the nodes and links in IDA's slipnet form perceptual symbol representations that carry forward throughout the entire architecture, including working memory, episodic memory (with a detour back through perception), long-term working memory, "consciousness," and action selection. There are no amodal representations.
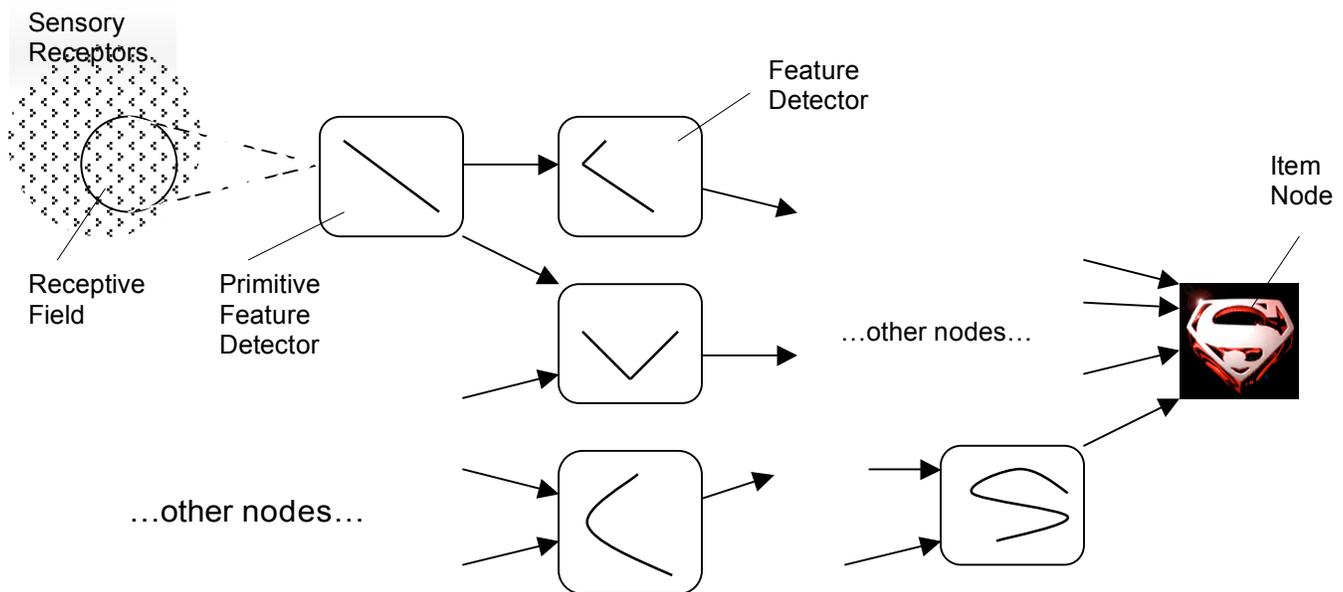

Figure 1. An early portion of a Slipnet. Primitive feature detectors respond to activity in their receptive fields and activate more complex feature detectors, etc. Feature detectors eventually combine to activate item nodes which, in turn, go on to activate category nodes. The process continues activating nodes of greater conceptual depth. When the Slipnet stabilizes, those nodes with total activation, that is combining base-level and current activation, above threshold, constitute the current percept. Thus, perception is a filtering process through which some nodes are selected to trigger local associations and to compete for conscious awareness.

Conclusion

What we've outlined above can be viewed as a program for implementing procedural memory and learning in autonomous software agents and robots. Such a program supports arguments for a developmental period, one of rapid perceptual learning, in the "lives" of both software agents and robots. Such a developmental period would circumvent the necessity of designing and implementing a complex ontology on the front end, a clear pragmatic advantage. In complex, dynamic environments, the learned ontology can be expected to out perform one designed and built in, and to do so with much less human effort.

Acknowledgements

References

Anwar, A., and S. Franklin. 2003. Sparse Distributed Memory for "Conscious" Software Agents. *Cognitive Systems Research* 4:339-354.

Baars, B. J. 1988. *A Cognitive Theory of Consciousness*. Cambridge: Cambridge University Press.

Baars, B. J. 2002. The conscious access hypothesis: origins and recent evidence. *Trends in Cognitive Science* 6:47–52.

Baars, B. J., and S. Franklin. 2003. How conscious experience and working memory interact. *Trends in Cognitive Science* 7:166–172.

Baddeley, A. D. 2000. The episodic buffer: a new component of working memory? *Trends in Cognitive Science* 4:417-423.

Barsalou, L. W. 1999. Perceptual symbol systems. *Behavioral and Brain Sciences* 22:577-609.

Bitterman, M. E. 1965. Phyletic difference in learning. *American Psychologist* 20:396–410.

Bogner, M., U. Ramamurthy, and S. Franklin. 2000. Consciousness" and Conceptual Learning in a Socially Situated Agent. In Human Cognition and Social Agent Technology, ed. K. Dautenhahn. Amsterdam: John Benjamins.

Edelman, G. M. 1987. *Neural Darwinism*. New York: Basic Books.

Edelman, G. M., and G. Tononi. 2000. *A Universe of Consciousness*. New York: Basic Books.

Ericsson, K. A., and W. Kintsch. 1995. Long-term working memory. *Psychological Review* 102:211-245.

Fahle, M., and I. Daum. 2002. Perceptual learning in amnesia. *Neuropsychologia* 40:1167–1172.

Franklin, S. 2000. Modeling Consciousness and Cognition in Software Agents. In *Proceedings of the Third International Conference on Cognitive Modeling, Groeningen, NL, March 2000*, ed. N. Taatgen. Veenendal, NL: Universal Press.

Franklin, S. 2000a. Deliberation and Voluntary Action in 'Conscious' Software Agents. *Neural Network World* 10:505-521.

Franklin, S. 2001a. Automating Human Information Agents. In *Practical Applications of Intelligent Agents*, ed. Z. Chen, and L. C. Jain. Berlin: Springer-Verlag.

Franklin, S. 2001b. Conscious Software: A Computational View of Mind. In *Soft Computing Agents: New Trends for Designing Autonomous Systems*, ed. V. Loia, and S. Sessa. Berlin: Springer (Physica-Verlag).

Franklin, S. 2003. IDA: A Conscious Artifact? *Journal of Consciousness Studies* 10:47-66.

Franklin, S., and A. C. Graesser. 1997. Is it an Agent, or just a Program?: A Taxonomy for Autonomous Agents. In *Intelligent Agents III*. Berlin: Springer Verlag.

Franklin, S., B. J. Baars, U. Ramamurthy, and M. Ventura. in review. The Role of Consciousness in Memory.

Franklin, S., and L. McCaulley. Feelings and Emotions as Motivators and Learning Facilitators. (Architectures for Modeling Emotions. AAAI Spring Symposia Technical Series Technical Reports; SS-04-02)

Gabrieli, J. D., W. Milberg, M. M. Keane, and S. Corkin. 1990. Intact priming of patterns despite impaired memory. *Neuropsychologia* 28:417–427.

Gould, J. L., and C. G. Gould. 1988. *The Honey Bee*. New York: W. H. Freeman.

Herrnstein, R. J. 1984. Objects, catagories, and discriminative stimuli. In *Animal Cognition*, ed. H. L. Roitblat, T. G. Beaver, and H. S. Terrace. Hillside, N.J.: Lawrence Erlbaum Associates.

Hobson, J. A. L. L. A. N., and R. O. B. E. R. T. Stickgold. 1994. A neurocognitive approach to dreaming. *Consciousness and Cognition* 3:16–29.

Hofstadter, D. R., and M. Mitchell. 1994. The Copycat Project: A model of mental fluidity and analogy-making. In *Advances in connectionist and neural computation theory, Vol. 2: logical connections*, ed. K. J. Holyoak, and J. A. Barnden. Norwood N.J.: Ablex.

Jackson, J. V. 1987. Idea for a Mind. *Siggart* Newsletter, 181:23-26.

Johnston, V. S. 1999. *Why We Feel:The Science of Human Emotions*. Reading MA: Perseus Books.

Kanerva, P. 1988. *Sparse Distributed Memory*. Cambridge MA: The MIT Press.

Kelemen, A., Y. Liang, and S. Franklin. 2002. A Comparative Study of Different Machine Learning Approaches for Decision Making. In *Recent Advances in Simulation, Computational Methods and Soft Computing*, ed. E. Mastorakis. Piraeus, Greece: WSEAS Press.

Maes, P. 1989. How to do the right thing. *Connection Science* 1:291-323.

Mandler, J. 2000. Perceptual and Conceptual Processes in Infancy. *Journal of Cognition and Development* 1:3–36.

Negatu, A., and S. Franklin. 2002. An action selection mechanism for 'conscious' software agents. Cognitive Science Quarterly 2:363-386.

Olton, D. S., J. T. Becker, and G. H. Handelman. 1979. Hippocampus, space and memory. *Behavioral and Brain Sciences* 2:313–365.

Pepperberg, L. 1990. Cognition in an African Grey Parrot. *Journal of Comparative Physiology* 104:41-52.

Yerkes, R. M., and J. D. Dodson. 1908. The Relationship of Strength of Stimulus to Rapidity of Habit Formation. *Journal of Comparative Neurology and Psychology* 18:459–482.

Zhang, Z., S. Franklin, B. Olde, Y. Wan, and A. Graesser. 1998. Natural Language Sensing for Autonomous Agents. In *Proceedings of IEEE International Joint Symposia on Intellgence Systems 98*.