Epistemic Structure: How Agents Change the World for Cognitive Congeniality[1]

Sanjay Chandrasekharan (schandra@sce.carleton.ca)
Cognitive Science Ph.D. Program, Carleton University,
Ottawa, Canada, K1S 5B6

Abstract

A two-part model is presented to explain how some organisms and humans add structure to the world to reduce cognitive complexity. Some philosophical possibilities of this model are then outlined, which include a situated cognition explanation of representation and intentionality.

Many organisms generate stable structures in the world to reduce cognitive complexity (minimize search or inference), for themselves, for others, or both. Wood mice (*Apodemus sylvaticus*) distribute small objects, such as leaves or twigs, as points of reference while foraging. They do this even under laboratory conditions, using plastic discs. Such "way-marking" diminish the likelihood of losing interesting locations during foraging (Stopka & MacDonald, 2003). Red foxes (*Vulpes vulpes*) use urine to mark food caches they have emptied. This marking acts as a memory aid and helps them avoid unnecessary search (Henry, 1977, reported in Stopka & MacDonald, 2003). The male bower bird builds colorful bowers (nest-like structures), which are used by females to make mating decisions (Zahavi & Zahavi, 1997). Ants drop pheromones to trace a path to a food source. Many mammals mark up their territories.

At the most basic level, cells in the immune system use antibodies that bind to attacking microbes, thereby "marking" them. Macrophages use this "marking" to identify and destroy invading microbes. Bacterial colonies use a strategy called "quorum sensing" to know that they have reached critical mass, (to attack, to emit light etc.). This strategy involves individual bacteria secreting molecules known as auto-inducers into the environment. The auto-inducers accumulate in the environment, and when it reaches a threshold, the colony moves into action (Silberman, 2003).

Such 'doping' of the world is commonly seen in lower animals, most large animals (large body & brain size) do not exploit this strategy. Except humans. More than any other species, humans generate external structure to reduce cognitive complexity, for themselves and/or for other humans. Markers, color-codes, page numbers, credit-ratings, badges, shelf-talkers, speed bugs, road signs, post-it notes, the list is almost endless. Humans also create external structure for reducing the cognitive complexity of artifacts. Examples include bar codes (help check-out machines' decisions easier), content-based tags in web pages (makes Web agents' decisions easier), sensors on roads (makes the traffic light's decisions easier), etc.

The pervasiveness of such structures across species indicates that adding structure to the world is a fundamental cognitive strategy (Kirsh, 1996). Note that these structures predominantly serve a task-smoothening function – they make tasks easier for agents. Some of these structures have referential properties, but they do not exist for the purpose of reference. They exist to make tasks easier. From here onwards I will term such stable structures for "cognitive congeniality" (Kirsh, 1996), *epistemic structures*. The term is derived from a distinction between epistemic and pragmatic action made by Kirsh (1994).

How do organisms generate and use such structures? This is the question I want to address in this paper. I am primarily interested in the human case, but I will explore the case of humans adding structure to the world in the broader context of other organisms exhibiting such behavior.

A Taxonomy and a Property

Most of the literature on epistemic structures is by David Kirsh, and from the Distributed Cognition tradition in general. Kirsh's work explores the structural and computational properties of such structures, and how they work. I am interested in the other half of the problem, i.e. how such structures are generated and used. I will use Kirsh's model of how such structures work to develop a model of how such structures are generated. Some philosophical possibilities presented by the model are then outlined, along with future work.

Epistemic structures can be classified into three, based on whom they are generated for. Examples in brackets.

1. Structures generated for oneself (Cache marking, bookmarks)
2. Structures generated for oneself and others (Pheromones, color codes)
3. Structures generated exclusively for others (Warning smells, badges)

A central feature of such structures is their task-specificity (more broadly, function/goal-orientedness). Here's an example to illustrate the concept. Think of a major soccer match in a large city, and thousands of fans arriving in the city to watch. The organizers put up large soccer balls in

[1] Carleton University Cognitive Science Technical Report 2004-03
URL: http://www.carleton.ca/iis/TechReports

every junction, and on streets leading up to the venue. Obviously, the ball reduces the fans' cognitive load, but how? To see how, we have to examine the condition where big soccer balls don't exist to guide the fans.

Imagine a soccer fan walking from his hotel to the game venue. She makes iterated queries to the world to find out her world state (What street is this? Which direction am I going?), and then does some internal processing on the information gained through the queries. After every few set of iterated queries and internal processing, she updates her world state and mental state, and this continues until she reaches her destination.

What changes when the ball is put up? The existence of the big soccer ball cuts out the iterated queries and internal processing. These are replaced by a single query for the ball, and its confirmation. The agent just queries for the ball, and once a confirmation comes in, updates its world state and internal state. The ball allows the agent to perform in a reactive, or almost-reactive mode, i.e. move from perception to action directly.

This happens because the ball is a *task-specific* structure -- it exists to direct soccer fans to the game venue. Other structures, like street names and landmarks in a city, are function-neutral or *task-neutral* structures. The fans have to take these task-neutral structures and synthesize them to get the task-specific output they want. Once the huge ball, a task-specific structure, exists in the world, they can use this structure directly, and cut out all the synthesizing. (How the soccer fans manage to discover the ball's task-specificity is a separate and relevant issue, but we will not address it here.) Task-specificity is a common property of all epistemic structures, including pheromones and markers.

Kirsh's model of "changing the world instead of oneself" (Kirsh, 1996), postulates that generation of external structures involve task-external actions, and the structures work by deforming the state space in a way that paths in a task environment are shortened. Such structures also allow new paths to be formed. Kirsh's model only tackles physical structures generated by organisms, like tools. He does not consider structures generated for cognitive congeniality. In the above example, I treat the task environment as a combination of the agent and the environment (what I term *action-environment*), a notion Kirsh resists. He thinks such a view devalues the generality of the notion of task environment. I argue below that this need not be the case.

The Tiredness Model

How are task-specific structures generated? We will first consider the case of organisms like ants, wood mice and red foxes. I will make two reasonable assumptions here. One, organisms generate random structures in the environment (pheromones, urine, leaf piles) as part of their everyday activity. Two, organisms can track their physical or cognitive effort (i.e. they get "tired"), and they have a bias to reduce cognitive effort or tiredness.

Now, some of the random structures are encountered while executing tasks like foraging and cache retrieval. In some random cases, these structures make the task easier for the organisms (following pheromones reduce search, avoiding urine makes cache retrieval faster, avoiding leaf-piles reduce foraging), i.e. they shorten paths in the task environment. Given the bias to avoid tiredness, these paths get preference, and they are reinforced. Since more structure generation leads to more of these paths, structure generation behavior is also reinforced. We have implemented two simulations to test this model, one using genetic algorithms and the other using Q-learning (See Chandrasekharan and Stewart, 2004). They show that a simple feedback of tiredness can lead to agents systematically using and generating external structures, both across generations and within their lifetime.

The tiredness model explains generation of structure in cases 1 and 2 (structures for oneself and structures for oneself & others) for organisms other than humans. Case 2 (structures generated for oneself & others) is explained by appealing to the similarity of systems – if a structure provides congeniality for me, it will provide congeniality for other systems like me. This is similar to how paths are formed in fields: one person cuts across the field to reduce his physical effort. Others, sharing the same system and wanting to reduce their effort, find the route optimal. As more people follow the route, a stable path is formed. This kind of self-organization, based on shared systems, is the reason why the combination of agent and environment does not violate the generality of the task environment notion.

For case 3, (structures generated exclusively for others), the model explains only some cases. For instance, it can explain the generation of warning smells and colors exclusively for others, because the effect of such structures can be formulated in terms of tiredness. (Release of some chemical cautions predators, this lowers the organism's fleeing response, thus reducing tiredness, which, when fed back, reinforces the generation system). However, this model cannot explain the generation of structures like the bower, and the peacock's tail, which do not seem to provide any tiredness benefit for the generator.

The same model, with some variations, works for humans as well. The variations are:

1. There is explicit awareness of tiredness, i.e. of harder paths in the task environment.
2. Structures are actively generated.
3. A reactive mode bias.

Once again, we have a task environment, with paths with different cognitive loads, or tiredness. Over several iterations of a task, the longer sections of a path (i.e. the ones with more cognitive load) become apparent to humans, and structures are generated specifically to shorten those sections. In the case of other organisms, structures are generated inadvertently (as part of their activity), and then discovered as reducing cognitive load. In the case of humans, structures are actively generated to minimize cognitive load.

How do generated structures get the nice property of task-specificity? This is where the reactive-mode bias comes into play. The reactive-mode bias seeks to minimize all paths in a task graph, i.e. "collapse" all longer paths to short perception-action sequences. Once a long (i.e. tiring) section of a path is identified in an action-environment (say involving search or long-winded inference), the reactive-mode rule seeks to "collapse" that section. This can be done using many techniques, like chunking (streamlining processes so that many disparate processes is replaced by one smooth process), delegation (hive processes off to another agent), or by pushing information about states to the world (so that the world does memory's job). Generating epistemic structure is a variation of the last technique. But instead of just states, an entire process is pushed to the world. This is done by making the output of the process exist in the world, so that it can be perceived, instead of computed. Doing this is simple: all you need is a rule that generates the output of the long path, and "stick it" at the point of perception.

Here's an example to illustrate the notion. Take the case of the epistemic structure known as shelf-talkers. These are little labels you find on the shelf in non-computerized stores. These labels usually contain redundant information, like the name of the product right above it. These labels are not for the benefit of shoppers, they exist to help store managers and clerks. If a store has thousands of products, and a product runs out, the store manager or clerk will have a hard time figuring out what was on the empty slot. These labels indicate to the clerk or manager what product was on the empty slot. If we model the store-manager's task in this situation, a long path in her action-environment starts from her discovery of the empty slot, and leads to her search for the identity of the product. Generating the shelf-talker involves taking the output of this search (the identity of the product), and sticking it to the starting point (the perception point) of the lengthy process, namely the empty slot. This provides a perception-action sequence right away. The structure is fine-tuned (say, size and manufacturer details added) over later iterations.

Four conclusions follow from this "path-collapse" model of epistemic structure generation. One, the task-specificity of epistemic structures is a direct outcome of the collapsing process that generates such structures. The task-specificity of a generated structure indicates the extent of identification and collapse of long paths by the generation algorithm. Two, in this model, structure can be generated only after at least one, usually many, iterations of a task. So epistemic structure generation is an indication of task expertise. Three, structures are generated on the basis of a cost trade-off – the cognitive benefit accrued from the generation of structures should be significantly higher than the (cognitive and physical) cost of generation of the structure, otherwise structures will not be generated. Finally, structures evolve in two ways. One, they are fine-tuned, as in the case of adding details to shelf-talkers. Two, they evolve as the agent's tasks change (say expiry date added to the shelf-talker for the new task of inventory management). Both these changes are accounted for by the reactive-mode bias.

The generation of structures for oneself and others is similar to the organism case and can be explained using the path-formation mechanism. A structure for oneself and others (like the shelf-talker) is actually a structure for oneself that others find useful because of shared systems.

The Curious Case of Others

What about case 3, structures generated exclusively for others, like badges and the soccer ball? Tiredness cannot drive the generation of structures here, because the generating agent is not executing the task. However, the only way to generate task-specific structures is to be aware of their cognitive load, i.e. by "running" them on your system. So how do we manage to generate task-specific structures exclusively for others, even though we are not doing the task? The answer is a two-part one. The first part says we are not very good at generating task-specific structures exclusively for others, especially for complex tasks. We get task-specificity through usability testing, by running the structures using others' systems (or one's own). But to get structures that can be usability tested, we need structures that at least approximate task-specificity. How do we manage to do this? Answer: by simulating others.

The notion of simulation used here combines three other notions of simulation, two existing and one new. The first notion is the one used in the propositional attitude literature (Nichols, Stich, Leslie & Klein, 1996), where simulation explains how we come to have beliefs about others. Here simulation is used primarily to understand the internal state of another agent. The second is the simulation heuristic, put forward by Kahneman and Tversky (1982) in counterfactual reasoning, to explain how participants "mutate" (change) different points in scenarios to generate alternatives to reality. Here simulation is used primarily to explain how we model the dynamic and temporal aspects of situations, and "mutate" points within it. The first notion of simulation is a requirement for the generation of case 3 epistemic structure, because structures are generated for another agent. The second notion of simulation is needed because epistemic structure generation requires modeling the dynamic and temporal aspects of a task, and generating alternatives to reality at different points, so that the task becomes easier. Finally, the notion of simulation here also incorporates using one's own system as a proxy for other systems, to "test-run" the efficiency of a generated structure.

We have run an experiment to test the simulation hypothesis in the case of structures generated exclusively for others (see Chandrasekharan, 2004). Participants were given three standard problem scenarios, faced by agents with different cognitive capacities (people from a hypothetical culture, blind people, Martians, robots etc.) They were asked to provide solutions to help these agents. Epistemic structures were generated, and they were task-specific, for agents cognitively close to the participants, but performance fell significantly on both counts as the cognitive distance between the participants and the agent increased. Participants were asked for self-reports on how they tried to solve the problems. Most of them claimed to simulate across all conditions, making simulation a necessary condition for generating task-specific structures, but not a sufficient one.

Philosophical Possibilities

The tiredness model has an interesting property: it combines teleological and processing approaches to cognition. Teleological explanations are usually avoided in cognition, though there is nothing wrong with them in principle. Teleological explanations differ from causal explanations in the following way (Papineau, 1995). A causal explanation explains a phenomenon by appealing to something that happened before it, i.e. a cause. For example, the ball moved because I kicked it. Teleological explanations explain a phenomenon by appealing to a consequence, i.e. something that happens after the phenomenon. For instance, some snakes developed venomous fangs because the venomous fangs help the snakes survive. This kind of 'forward-looking' explanation is given in areas involving living entities, like in biology and economics.

I do not appeal to natural selection or genetic transmission directly while trying to give my aetiological account. Instead, I appeal to an intermediary level, a processing level that exists below functions, but above natural selection. Specifically, I appeal to computational load or effort that underlie cognitive events. The basic idea is that organisms can sense cognitive exertion (or computational load), just like physical exertion. And just as organisms develop features that help them minimize physical exertion (claws, tusks, etc.), they can develop features that let them minimize cognitive exertion. Essentially, cognitive load and its internal tracking lead up to epistemic structures. Processing models are not usually used this way (as an aetiological account), and this provides some interesting possibilities, a major one being a unified account of cognition that brings together computational and biological explanations.

Processing load by itself does not provide a full aetiological account, because processing effort bottoms out to energy conservation and fitness. So the bottom level would be natural selection, driven by the advantage provided by energy conservation (Our genetic algorithm simulation is based on this idea.) The processing level is hidden in current aetiological accounts, I just abstract it out as a separate level.

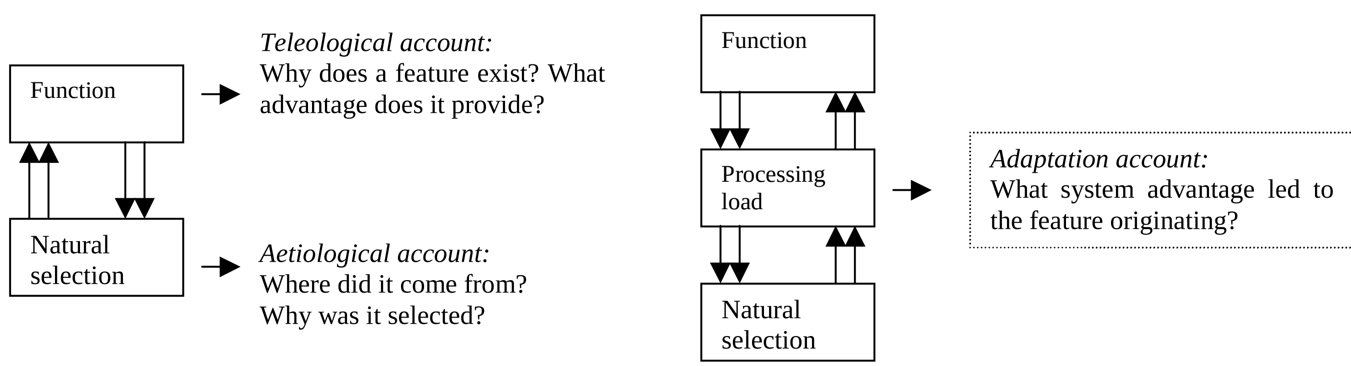

Figure 1: The tri-level model applied to biological explanations of cognition

Teleological explanations usually have a second step, known as the aetiological account, which explains how the feature originated and came to be where we find it (Bogan, 1995). Two kinds of mechanisms are usually considered to explain how a particular feature came about. One is genetic transmission, whereby features are passed from one generation to the next. The other is selection mechanisms (better energy management, for instance), whereby organisms with a particular feature have a better chance to reproduce than organisms which lack it.

The idea behind this aetiological explanation is that it is not enough to appeal to functions, we should have an account of underlying material processes that operate, leading to a feature existing. The teleological account explains *why* a feature exists, the aetiological account explains *where* it came from, and *what* led to it being selected.

The processing level is considered the adaptation level, where variations in system variables (processing load, memory etc) lead an individual organism to develop a cognitive feature. Introducing this intermediary level in a teleological account has the (satisfying) effect of combining Marr's tri-level model with situated and biological explanations of cognition. In figure 1, the model on the left shows teleological explanations without the processing level, and the model on the right with the processing level added.

What advantages does this new level offer? One, it allows us to move away from the messy details of prehistoric environments and selection processes, and work with a more manageable computational account. Two, if we just consider computational load or effort, it provides us with a physical mechanism that all kinds of organisms can possibly monitor, without explicit awareness of the process. In fact, event-related fMRI studies have shown that automatic on-line monitoring of response and task-difficulty happens in humans, so that higher levels of control and attention can be deployed to avoid erroneous responses (Hopfinger et al, 2001). This kind of automatic monitoring of computational load provides us with a system variable common to a species, and can explain how some cognitive adaptations originated in that species.

Besides this, the model offers a possibility to explain the nature of representation. Our current simulations implement an evolutionary process that leads to organisms generating task-specific external structures in the world. These structures lower cognitive load and are accessed at run-time, while organisms execute tasks. Interestingly, the same model can explain generation and tracking of internal structures in organisms.

To see how, consider foraging bees. Assume that some sequence of memory traces of landmarks (say a tall tree, a lake, a garden) are left in their brain as a result of their everyday foraging activity. In some foraging trips of some bees, the trace sequence match, to some degree, the external structures they perceive (tall tree → lake → garden). Such trips involve less search, because they lead to food directly, i.e. they form shorter paths in the task environment. Over time, the bias against tiredness leads to such paths being used more, they are reinforced. This leads to landmark-based navigation, which exist in bees (Gould, 1990). As in the case of external structures, the generation of memory traces is reinforced because more traces lead to more such shorter paths in the task environment. We are currently working on simulating this example.

This model presents a situated cognition explanation of how memory structures come to be used as task-specific structures, and why such internal structures are systematically generated. If such task-specific memory structures are considered to be representations (that is, as standing for something specific in the world), then the model explains, in a computationally tractable manner, how representations are generated and used.

The model also explains what such ‘primitive’ representations are -- they are internal traces of the world that allow the agent to shorten paths in a task environment. Roughly, they are computation- reducing structures (and equivalently, energy-saving structures). They are internal “stepping stones” that allow organisms to efficiently negotiate the ocean of stimuli they encounter. This means the traditional cognitive science view, that thinking is computations happening over representations, presents a secondary process – it describes a privileged path in the task environment. In the stepping stone view, representations are crucial for organisms, but they are just useful and incidental entities, not fundamental entities.

Moving up a level, the model could be used to explain intentionality, though not fully, given the model’s teleological roots. However, if we break down intentionality into two levels, Intentionality 1 and Intentionality 2, the model can explain the latter to some extent. Intentionality 1 is the basic directedness every living organism has towards food, mates and survival, it is the directedness shared by humans and bacteria. This is the problem of life, and it is a problem in molecular biology. Intentionality 2 is higher-order directedness, the directedness towards secondary structures, i.e. structures that are only indirectly “about” food, mates or survival. The tiredness model explains how directedness towards such secondary structures can come about, and how this leads to the systematic generation of such internal structure by organisms. Since it is based on a simple feedback mechanism that tracks energy levels, the model provides a rudimentary physicalist explanation for second-order directedness -- it is a computation-reducing/energy-preserving mechanism.

This obviously does not provide a model of how intentionality is implemented in the brain, but it could be used as a starting point to develop such an explanation.

Acknowledgment

Some of the ideas presented in this paper were developed while I was a pre-doctoral fellow with the Adaptive Behavior and Cognition (ABC) Group of the Max Planck Institute for Human Development, Berlin. I would like to acknowledge the group's support, particularly encouragement and critical feedback from Dr. Peter Todd and Dr. John Hutchinson.

References

Alcock, J. (1998). *Animal Behavior: An evolutionary approach*, Sunderland, Mass., Sinauer Associates.

Bogen, J. (1995). Teleological explanation. In Honderich (Ed.), *The Oxford Companion to Philosophy*. New York: Oxford University Press

Bradbury, J.W. & Vehrencamp, S.L. (1998). *Principles of Animal Communication.* Sunderland, Mass: Sinauer Associates.

Chandrasekharan, S. & Stewart, T. (2004). Reactive agents learn to add epistemic structures to the world. Carleton University Cognitive Science Technical Report 2004-01 http://www.carleton.ca/iis/TechReports/

Chandrasekharan, S. (2004). Solutions that change the world are not readily available. Carleton University Cognitive Science Technical Report 2004-02 http://www.carleton.ca/iis/TechReports/

Clark A. (1997). *Being There: putting brain, body, and world together again*, Cambridge, Mass., MIT Press.

Gould, J.L. (1990) Honey bee cognition. *Cognition*, 37, 83-103.

Henry, J.D. (1977). The use of urine marking in the scavenging behaviour of the red fox (Vulpes vulpes). *Behaviour*, 62:82-105.

Hopfinger, J. B., Woldorff, M. G., Fletcher, E., & Mangun, G. R. (2001). Dissociating top-down attentional control from selective perception and action. *Neuropsychologia*, 39, 1277-1291.

Kahneman, D., & Tversky, A (1982). The simulation heuristic. In D. Kahneman, P. Slovic, & A. Tversky, (Eds.), *Judgment under uncertainty: Heuristics and biases*. New York: Cambridge University Press.

Kirsh, D., & Maglio, P. (1994). On distinguishing epistemic from pragmatic action. *Cognitive Science*, 18, 513-549.

Kirsh, D. (1996). Adapting the environment instead of oneself. *Adaptive Behavior*, Vol 4, No. 3/4, 415-452.

Mandik, P. and Clark, A. (2002). Selective Representing and World Making. *Minds and Machines* 12: 383-395.

Millikan, R. G. (1994). Biosemantics. In Stich, S. and Warfield, T. A. (Eds.), *Mental representation, a reader*. Oxford: Blackwell.

Millikan, R.G. (1996). "Pushmi-pullyu Representations", in James Tomberlin, ed., Philosophical Perspectives vol. IX, (Ridgeview Publishing) pp. 185-200. Reprinted in *Mind and Morals*, ed. L. May and M. Friedman, pp. 145-161. Cambridge, Mass. MIT Press

Nichols, S., Stich, S., Leslie, A., Klein, D. (1996). Varieties of off-line simulation. In: Carruthers D and Smith E (Eds.) *Theories of Theories of Mind*, Cambridge: Cambridge University Press.

Papineau, D. (1995). Biology, philosophical problems of. In Honderich (Ed.), *The Oxford Companion to Philosophy*. New York: Oxford University Press.

Silberman, S. (2003), The Bacteria Whisperer. *Wired*, Issue 11.04, April 2003.

Stopka, P. & Macdonald, D. W. (2003) Way-marking behaviour: an aid to spatial navigation in the wood mouse (Apodemus sylvaticus). *BMC Ecology*, published online, http://www.biomedcentral.com/1472-6785/3/3

Zahavi, A., & Zahavi, A. (1997). *The Handicap Principle: A missing piece of Darwin's puzzle*. Oxford: Oxford University Press.