

A Distributed Cognition Perspective on Symbiotic Cognitive Systems: External Representations as a Medium for Symbiosis

Thomas Erickson

IBM T. J. Watson Research Center
snowfall@us.ibm.com

Abstract

This paper offers a perspective on Symbiotic Cognitive Systems that draws on Distributed Cognition. It argues that representations are the medium of cognition, and that the external representations that are one of the foci of Distributed Cognition are critical to supporting symbiosis. The paper analyzes an instance of a symbiotic cognitive system in which hundreds of human participants – with the support of a digital system – collectively optimize a program. It discusses the roles external representations play in symbiosis, and suggest that the design of external representations that are accessible and legible to both human and digital agents is a critical part of symbiotic cognitive systems.

What Does it Mean to be Cognitive?

What does it mean for something to be “cognitive?” In my view, there is no consensus. Cynics might say that “cognitive” is simply AI re-branded, perhaps with sensors or big data thrown in. A more principled view is that “cognitive” means a system that possesses macro-scale properties – perception, reasoning, learning/memory and action – that characterize systems (i.e. humans) that all agree are cognitive. Taking this position, one might develop a taxonomy of systems: the simplest being a sense-analyze-respond system, and more complex version having added features like the ability to alter goals, move about their environment, or cooperate with others. Different people will draw the ‘cognitive line’ at different levels of this taxonomy.

In this paper I argue that regardless of where we draw the cognitive line, the design of representations to enable cognition is central. Examining the ways representations support cognitive activity will yield actionable implications for analyzing and designing cognitive systems.

My starting point is the theory of Distributed Cognition, as articulated by Hutchins (1995). Its core insight is that cognition does not occur solely within an individual mind,

but rather is distributed across people, artifacts and environments. While cognition can occur solely in a person’s head, it is more common for it to spill out into the world. Distributed Cognition argues that people think by creating and modifying representations, often *external* representations. “External” means that the representation exists outside the mind in a material form such as a list or diagram.

The external aspect of representations is critical: if a representation is external, it means that it can be accessed and changed by multiple agents (human or digital), and that cognition is a process that can occur in groups and systems. This paper will (1) discuss external representations and how they support cognition; and (2) provide an example of how external representations function in a symbiotic cognitive system. The paper will argue that for symbiosis to succeed, it is important that external representations be accessible and legible to both human and digital agents.

External Representations and Cognition

Lists

A simple example of an external representation is a list. Figure 1 shows two versions of a list of medications prescribed to a man who had recently had a stroke (this is from the work Ruth Day, cited in Norman, 1993).

The image on the left is the form in which the prescription list was given to the patient. This list is arranged to suits the needs of physicians and pharmacists – it makes it easy to look for any medication and see the frequency and

Inderal	— 1 tablet 3 times a day with meals				
Lanoxin	— 1 tablet at breakfast				
Carafate	— 1 tablet with meals and at bedtime				
Zantac	— 1 tablet at breakfast and at dinner				
Quingalute	— 1 tablet 4 times a day				
Coumadin	— 1 tablet at bedtime				

	Breakfast	Lunch	Dinner	Bedtime
Lanoxin	X			
Inderal	X	X	X	
Quinaglute	X	X	X	X
Carafate	X	X	X	X
Zantac		X		X
Coumadin				X

Figure 1. Isomorphs of a list that supports different tasks.

timing with which it was prescribed. However, a problem with the list on the left is that it is not very easy for an ordinary person (let alone a stroke victim) to use. It does not make it easy to answer practical questions like: “It’s breakfast time, what do I take now?” or “I’m leaving the house at 10am, and won’t be home until 9pm, how many of each pill should I bring with me?”

The image on the right (Figure 1) shows an isomorphic version of the list: the content is the same, but the layout was designed with the patient’s questions in mind: one can see, at a glance, which medications need to be taken when.

The point is not that one version of the list is better than another, but that the versions are designed to support different cognitive tasks. In a paper-based world it is not a bad idea to standardize on one form of representation; but in the digital world, where representations are mutable, it makes sense to think about supporting multiple versions of a representation so that it can easily support multiple tasks.

Other Representations

While, as representations go, lists are about as simple as one can get, the point is generally true: the design of representations impacts the ease of performing cognitive tasks. There is a nice set of examples of external representations in Don Norman’s book, *Things that Make Us Smart* (Norman, 1993). They range from numeric symbols (Arabic vs. Roman Numerals) to physical representations (the rings and pegs of the classic “Tower of Hanoi” puzzle). The lesson is that, using metrics such as performance and error frequencies, the design of representations matter. Two representations that are equivalent in their content may differ in the degree to which they facilitate or inhibit the performance of various cognitive tasks.

Symbiotic Cognition: An Example

Thus far the discussion has focused on representations that enable an individual to carry out a cognitive task. In this section, I discuss an example in which external representa-

tions are used to support symbiotic cognition. The example to be considered is an online “contest” in which hundreds of human agents and one (or more, depending on how one counts) digital agents, collectively optimize a program.

In late 1998 a company called The Mathworks began holding online programming contests to promote its programming language, MATLAB. The contest worked as follows. A problem was proposed – for instance, guide 5 Mars Surveyor robots in exploring a map of passable and impassable regions – and contestants would write programs to solve it. Upon submission, each contestant’s program was automatically scored, the score reflecting two factors: (1) quality – percent of territory explored by surveyors; and (2) performance – the CPU time required. Once scored, the results were posted on a *Standings Page* (Figure 2) that ranked all submissions by their scores, identified the authors, and disclosed both the quality and performance metrics that went into each score.

What makes this contest of particular interest is that it was ‘open source:’ once a submission was scored and posted on the *Standings Page*, a single click revealed the entry’s code. This enabled the practice of “tweaking.” For instance, contestant A might devise a new algorithm for solving a problem that vaulted her into first place, but a few minutes later contestant B might examine her code, replace a single statement with a more efficient version, and take over first place due to slightly better performance. Just as tweaking was quick and easy for B, so would it be for C, D and E, and thus B’s first place status was generally short lived. The *Standings Page* showed a lot of churn.

Tweaking was controversial, as Ned Gulley describes in his account of the contest (Gulley, 2005). Many participants were ambivalent – they did it, but felt bad about it. This is visible in the names of entries like “ripoff,” “outright theft,” and “If All Else Fails, TWEAK!” Others were less apologetic, with chiding names like: “Brackets are Expensive.” Others paid homage to their ancestors: an entry named Soup Dragon, whose descendants were tweaked into ascendancy for a time, generated offspring with names like “SoupMix,” “SneakyGreenSoup,” and the ultimate

Status: All			In Queue	Scored	1 - 100 of 1371		
Submit Date	Rank ▲	Title	Player	Results	CPU Time	Score	
Jun 18 16:51	1	NoSoup4U !1	Paul Uribe	4	4.677	109.282	
Jun 18 16:53	2	ddebfinal2	Debora Poon	4	4.667	110.724	
Jun 18 16:53	3	debfinal	Debora Poon	4	4.677	110.774	
Jun 18 16:53	4	debfinalbackup	Debora Poon	4	4.687	110.824	
Jun 18 16:51	5	sneak_a_little_opt2	JB	4	4.586	111.562	

Figure 2. The Standings List serves as an external representation that supports collective program optimization in the MatLab contest.

winner, “NoSoup4U!” As Gulley notes, in spite of the ambivalence, “*tweaking turns out to be the fuel that drives the entire contest. It offers an immediate reward to the tweaker – for an investment of a few minutes, your name appears on the leader board [Standings Page], wreathed in glory. The practice is also a call to arms for the original author (“How dare someone tweak my code!”). If you get tweaked, you want to know about it. And you may work very hard to tweak your way back into first place.*”

Not only does tweaking engage contestants, but collectively it is a powerful optimization mechanism. Figure 3 shows an analysis of the second contest, focused on the “Mars Surveyor” problem. The figure shows each submission as a dot, positioned as a function of its time stamp and score, so that the lowest dot at any given time designates the current first place entry. The dots connected by the line show the gradual optimization of an instance of an algorithm. A general pattern, within and between contests, was that promising new algorithms would be rapidly optimized in crescendos of tweaking.

This is an excellent illustration of a large scale cognitive system composed of human and digital agents. Let’s examine the role that external representations play.

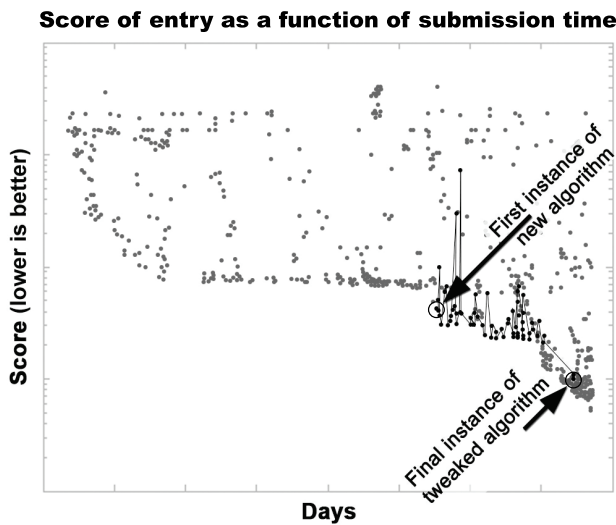


Figure 3. Collective optimization of contest entries through repeated ‘tweaks’ by multiple participants. The dots connected by the line show the appearance of a novel algorithm that is repeatedly tweaked. (The Mathworks, 2015)

The External Representations and their Roles

The cognitive system – that is, the MatLab contest comprised of its human agents and digital infrastructure – consists of three external representations that enable human agents to collectively optimize a program for solving a particular problem. It is these representations – to which both digital and human agents have access – that enable collaboration among individuals (even though the ‘collaboration,’ being part of a contest, is not entirely voluntary).

There are three external representations in play: (1) the *Rules Page*; (2) the *Submissions*; and (3) the *Standings List*. The *Rules Page* is quite simple: it describes the problem to be solved, provides resources such as sample code, test functions, a forum, and lays out the rules of the contest – including making it clear that ‘tweaking’ is encouraged. Basically, it provides what participants need to get started. The other two external representations – **the Submissions** and **the Standings List** – play more complex roles as described below.

The Set of Submissions Enables Integration

The *Submissions* – which we consider as a collective – is the growing cloud of contest entries that solve the same problem, many of which are tweaked versions of previous submissions. The *Submissions*, in their entirety, capture elements of the solution, and optimizations of its components. Because participants can download any submission, they can analyze and modify it as they wish, and then upload the new version.

What is important is that this external representation enables tweaks to accumulate: A’s submission has a novel algorithm; B downloads and tweaks A’s code to make it faster and resubmits it; C downloads and tweaks B’s code and resubmits that, and so on...

But why does B decide to tweak A’s submission in particular? There may have been 100 submissions since B last visited. And why did C focus on B’s code rather than any of hundreds of other submissions?

The Standings List Enables Cognitive Focus

A critical feature of the contest is that the cloud of program instances is not simply used by human agents, it is accessible to the digital elements of the system. That’s important because, loosely speaking, the system understands the *Submissions*: it is able to run each submission and give it a score, and position it in the *Standings List*. And it is the *Standings List* that enables participants to decide where to focus their attention in the *Submissions*. Most attention will likely be focused on the current leader, but it is also the case that certain authors (whose names are listed in the *Standings List*) or certain approaches (which are both associated with authors and often signaled in the submissions’

names) may develop followings. Whatever the case, the *Standings List* enables B to find A's submission, and C to find B's. In short, *the Standings List* provides an index into *the ever-growing cloud of Submissions*.

The *Standings List* Enables Incentive Mechanisms

The Standings List plays a second role as well. It serves as an incentive mechanism. Because this large-scale computation is dependent on humans to function, the system needs to encourage participation. There are many approaches to this, but in this case the activity has been framed as a "contest." That means it has "winners" and "losers," that "competition" is expected, that the "contest" has a "start" and an "end," that there is often a burst of effort as the end nears, and on and on. It is this framing, and its explicit support via the *Standings List*, that provides the motive power for the collective interaction. And at the same time, it is this framing, and specifically contest-linked values having to do with fairness and cheating, that makes contestants ambivalent about tweaking, even though it is explicitly encouraged in the rules. The ability to design and operate the MatLab contest in a way that simultaneously tapped the incentive power of the contest framing even as it positioned tweaking as "OK" rather than cheating, is what makes this an artful example of symbiotic cognitive systems design.

Summary: A Symbiotic Cognitive System

In the MatLab contest, participants submit programs that solve a complex problem. Upon submission, programs are digitally evaluated for their speed and quality, and they (and their authors) are positioned appropriately in the public *Standings List*. Because submissions can be downloaded, modified and resubmitted by other participants, the contest encourages tweaks: small changes that make a program better and thus enable the tweaker to beat the original submitter in the *Standings List*. Multiplied across contestants and algorithms, this collective tweaking proves to be a powerful form of optimization.

I argue that the MatLab contest is a good example of a symbiotic cognitive system. Each instance of the contest is a 'run' of the system, and produces an optimized program that solves a complex problem. While this cognitive system is admittedly human-centric, note that it would be impossible without the digital evaluation and maintenance of its external representations, *the Set of All Submissions* and *the Standings List*.

Finally, note that other human-centric symbiotic systems are amendable to similar analyses, and offer more roles for digital agents (e.g., Wikipedia with its various bots for countering vandalism, tagging, and validation).

Conclusion and Proposal

This position paper argues that the design of external representations is central to symbiotic cognitive systems. First, it has noted that different but isomorphic forms of a representation can make a particular cognitive task easier or more difficult. Second, it has analyzed a large-scale distributed cognitive system in which hundreds of participants, without explicitly cooperating, collectively optimize programmatic solutions to complex problems. The analysis shows how two external representations – the *Set of All Submissions*, and the *Standings List* – enable a system that **incent**s individual agents to make 'tweaks' to a program, permits individuals to **focus** on the particular submissions that interest them so they may be tweaked, and thus allows **integration** of optimizations over time and across participants.

While there has been considerable focus on how representations – whether used by digital or human agents – can increase the efficiency with which a computational system can solve problems, less attention has been given to how external representations play other roles. In particular, with a focus on symbiosis, I think it could be useful for the workshop to try to think through some of the ways in which external representations might support symbiotic activity. The MatLab contest illustrates how external representations can support integration, attention, and incentives; doubtless there are other ways in which external representations can serve these ends. Furthermore, there are no doubt other activities necessary to symbiotic cognition – planning, ranking and selection, achievement of consensus, generation of variations on a concept, generalization, supporting cooperation among agents, impedance matching between human and digital agents, and so on – that external representations can support.

I am not sure of the best way for the workshop to proceed were we to pursue this path, but I think a focus on representation would allow all participants — regardless of where they draw the cognitive line, and what they view as a cognitive system – to contribute to a symbiotic result.

References

- Gulley, N. 2004. In Praise of Tweaking. *Interactions*, 11:3: 18–23.
- Hutchins, E. 1995. *Cognition in the Wild*. Cambridge, MA: MIT Press.
- Norman, D. 1993. *Things that Make Us Smart*. Reading, MA: Addison-Wesley.
- The Mathworks. The Mathworks Mars Surveyor Contest Blog: <http://www.mathworks.com/matlabcentral/contest/contests/4/statistics.html>. Accessed 21-October-2015.