

Offloading Cognition onto the Web

Leslie Carr & Stevan Harnad

School of Electronics and Computer Science

University of Southampton

Highfield, Southampton SO17 1BJ

<http://users.ecs.soton.ac.uk/harnad>

<http://users.ecs.soton.ac.uk/lac>

ABSTRACT: *In modeling human cognitive capacity there is a question of what needs to be built in and what can be left out, because we can offload it onto cognitive technology, such as a google web search. Word meanings can be represented in two different ways: sensorimotor and verbal. Verbals definitions and descriptions can be offloaded, sensorimotor representations cannot. Dictionaries have a "grounding kernel" of words from which all other words can be reached through recombinatory definition alone. The words are learned at an earlier age and are more concrete. We tested conjunctive and disjunctive google search for target terms that had their own wikipedia entries, using either the target terms themselves, or the three words that had the highest co-occurrence frequency (latent semantic analysis) with the target words in Wordnet. The highly co-occurring words were surprisingly ineffective in retrieving the target word, even in joint conjunctive and disjunctive searches and there was no significant correlation with age of acquisition or concreteness. This raises some questions about the similarity between human associative memory and google-based associative search.*

Implicit and Explicit Know-How. Neuropsychology and neuroimaging studies have confirmed what we all knew already from introspection: That some of our know-how is conscious, but most of it is not. Learning, skill, knowledge and memory all come in two forms: "explicit," in which we are aware of and can hence describe in words *how* we are able to do what we are doing, and "implicit," in which we can do what we do, but we have no idea how we are doing it. Most of cognitive science is devoted to explaining how we are able to do what we can do by trying to discover the implicit (unconscious) mechanisms underlying our cognitive competence and making them explicit. Conscious introspection does not reveal what they are. The explanatory goal of cognitive science is to reverse engineer what it takes to pass the Turing Test (TT): Once we can successfully design a system that is capable of doing whatever any person can do, indistinguishably *from* any person, *to* any person, then we have a candidate explanation of how the human brain does it.

What Know-How Cannot Be Offloaded Onto Cognitive Technology? Many variants of the TT have been mooted, most of them irrelevant to explaining cognition, but we will consider one thought-experimental variant here so as to illuminate certain cognitive questions: The TT is a comparison between human and machine, to determine whether or not the performance capacity of the machine is distinguishable from that of the human. In the online era, increasingly powerful cognitive technology is available for people to use to do what they formerly had to do in their heads, as well as to enhance their performance capacity. One could even say that it is becoming possible to *offload* more and more cognitive processing onto cognitive technology, liberating as well as augmenting the performance power of the human brain. What effect – if any – does this have on the search for the underlying mechanism that cognitive science is trying to model and the TT is trying to test? How much (and what) of cognitive capacity *cannot* be offloaded onto cognitive technology (Dror & Harnad 2009)? The question is related – but in a somewhat counterintuitive way -- to the difference between

what I can do that I *know how I do*, and what I can do *without knowing how I do it*. It is our implicit “know-how” (the know-how that we have without knowing how) that can be offloaded onto technology without our even noticing a difference. Yet it is the explicit know-how that we can verbalize and formalize that is the easiest to offload (because we know how we do it, so it is easier to turn it into an algorithm).

Unconscious Know-How. Suppose we are trying to recall someone’s name: We know the person. We know that we know their name. But we just can’t retrieve it. And then it pops up: We have no idea how, or from where. When the name pops up instantly, with no tip-of-the-tongue delay, we take it for granted. But in neither case -- instant or delayed -- do we have any idea *how* we managed to retrieve the name. The need for a mechanism is even more obvious when it is not just rote memory retrieval that is at issue, but active processing: the execution of an algorithm. Most of us know the multiplication table up to 12x12 by rote: Beyond that we have to perform a computation. When we perform the computation externally, with pen and paper, it is evident what computation we are doing, and we can describe that same computation explicitly even when we do it in our heads. But there are computations that are done for us implicitly by our brains, computations that we cannot verbalize, nor are we even aware of their algorithm or its execution. Catching a Frisbee is as good an example as any: A robot doing that has to do certain optokinetic computations; so do we. But we are not aware of doing them, nor of what they are, while they are happening, implicitly. The same is true when we are trying to recall the name of a face we have recognized: It is not just the name-retrieval process that is executed for us, implicitly, by our brains, but the process of recognizing the face: As attempts to generate the same know-how in robotic vision show, a good deal of computation is needed there too, computation that is again completely opaque to our conscious introspection.

In and Out of Mind. So we are raising a question here about the *locus* of all that unconscious data-storage and data-processing underlying our cognitive know-how: How much of it really needs to take place inside the brain? Some have suggested that when one finds an item by searching google instead of searching one’s brainware, google becomes a part of one’s “extended mind”: But does that mean we need to build google into our TT candidate, or just that the TT candidate as well as our human subject should have access to google? And is what we are not and cannot be conscious of really part of our mind? We don’t just offload cognitive processing onto technology, we also offload it onto one another’s brains: Cognition is not only distributed but collaborative. Yet the other minds with which I collaborate do not become a part of my own extended mind. Nor would we want to design a TT candidate that included the brainware of other people, any more than we would want to include the hardware, software or data of external technology that our candidate merely uses.

Sensorimotor Grounding. People differ in their cognitive skills, and the TT is only meant to capture our generic human cognitive capacity, the know-how we would expect anyone to have. The distance between a layman and what used to be reckoned a “scholar” has definitely been narrowed by google (and google scholar!), for example. A layman plus google can now answer questions that only a scholar could answer, after a lot of library research. And cognitive technology can even endow average minds with what used to be considered autistic savant skills, performing instant computations that used to take a lot of real time. We accordingly need to ask the question: What are the cognitive capacities that *cannot* be offloaded onto cognitive technology -- the ones for which google is of no help? The obvious candidate is basic sensorimotor skills: Google may be able to help you improve your golf skills, but you would not want to *learn* to play golf (or acquire any continuous sensorimotor skill) bottom-up from google, any more than from an instruction manual. The same is true of wine-tasting. Some things can only be learnt from direct sensorimotor experience. Language itself is already cognitive technology, and its enormous power is derived from its combinatorial capacity, just as google’s is. Once you have grounded the meanings of a certain number of words directly through sensorimotor category learning, then language, books, algorithms and google can go on to combine and recombine those grounded words in myriad useful ways. But their initial grounding always has to be internal and sensorimotor.

The External and Internal Lexicon. We have shown, for example, that any dictionary can be systematically reduced to a much smaller “grounding kernel” of words from which all the rest of the words in the dictionary can be “reached” through verbal definitions alone, using recombinations of the grounding words (Blondin Massé et al 2008). Overall, those grounding words turn out to have been learned at an earlier age, and to be more concrete and sensorimotor than the rest of the dictionary, based on the MRC psycholinguistic database (Wilson 1988); but when we partial out the effects of the strongest correlate – age -- then the residual grounding kernel words are more abstract than the rest of the dictionary (Harnad et al 2008). Further analyses are showing that the grounding kernel consists of a highly concrete core, learned early, plus a surrounding abstract layer that is not correlated with age; hence some of the grounding vocabulary is learned early and some of it later. We are hypothesizing that the mental lexicon encoded in our brains is homologous to our external lexicon, in that it too has a grounded kernel of mostly concrete words that we acquired early, nonverbally, plus a further grounding layer of more abstract words; the meanings of these directly grounded words are encoded in nonlinguistic, sensorimotor representations. The rest of our word meanings are encoded verbally, as combinations and recombinations of these grounding words, for the most part, although no doubt some new words continue to be grounded directly in experience throughout the life cycle rather than just being encoded as verbal definitions, descriptions and explanations; and perhaps even some of those words that are initially learned explicitly from verbal definitions eventually become grounded more directly in experience with time; or perhaps, with frequent linguistic use, they are “rechunked” into more implicit, holistic units that we would have almost as much difficulty defining explicitly as we would in defining our core grounding vocabulary itself.

Boolean Search on the Web as Reverse Definition (Jeopardy) Search. The difference between a dictionary and an encyclopedia is really just a matter of degree. We expect more detail from an encyclopedia, and we consult it even if we already know what a word means, if we want more details about its referent. Hence it is natural to extend to encyclopedias the inquiry into the implicit and explicit encoding of meaning, but encyclopedia entries contain too many words. In our dictionary analyses, we ignored syntax and treated all the words in a definition as an unordered string of words (ignoring function words). Nevertheless, dictionary definitions proved short enough so we could converge on the grounding kernel without inordinate amounts of computation. If we think of the Web as an encyclopedia, almanac and vademecum, then there are two ways we can consult it for information: One is conventional dictionary- or encyclopedia-style look-up on the defined target term, say, “conscience,” for the defining terms: “feeling of right and wrong”. The nature of dictionary and encyclopedia search and use is that we rarely want to go in the opposite direction: from the defining terms to the defined term. In searching the web, however, this is often the direction we need to go. This is rather like playing “Jeopardy,” the game of reverse definitions, where you are given *parts* of the definition and must guess what is being defined.

Wikipedia As Search Success Criterion. Google uses Boolean AND/OR/NOT search plus a lot of background heuristics. The result is that for many if not most single-word searches based directly on the defined target term, the top google hit will be the Wikipedia entry for that term. Such a search is successful by all criteria of success if the first hit is the direct Wikipedia entry when an exact Wikipedia entry exists for the target search term. Wikipedia is an encyclopedia, and one of the most highly used “authorities” on the web based on the PageRank algorithm (in terms of links as well as hits). But a direct hit using the target term itself is trivial. The search becomes more realistic and challenging if the search term is not just the target term itself, or not the target term alone, but rather something we think is somehow related semantically to what we are seeking (when we are seeking, but do not know, the target term). In our brain, assuming that the target term and its meaning is in there somewhere, this would be like cued memory retrieval, or the “tip of the tongue” phenomenon, where the word associated with the target word we are seeing somehow triggers an unconscious process that retrieves the target word.

Boolean “AND” Search. There is no doubt that our “desert island” search function (if we were allowed only one) would be *conjunction*: Suppose you have a partial descriptor, S1, that is not specific

enough to retrieve what you are looking for by making it appear at or near the top of the resulting Google list of hits. Then if you have a second partial descriptor, *S2*, a *disjunctive* search (*S1 OR S2*) widens the search's extension still more, further reducing the chances that the target will come up near the top of the list, whereas a *conjunctive* search (*S1 AND S2*) adds the google extension of *S2*, but narrows the joint outcome to only the overlap between the extensions of both search terms. As a first approximation, a good choice of *S2* will greatly increase the chances of retrieving the target with a high rank. In fact, conjunction is so powerful that it is one of those formal properties – alongside the law of the excluded middle and the invariants of projective geometry -- that the brain would be downright foolish not to exploit. So it is likely that conjunction-based retrieval of the joint-associates of two or more pairs of “search terms” plays a role in the cognitive process that underlies trying to recall or even to generate concepts. If so, then associative retrieval based on conjunctive search might be one of the kinds of process that can be offloaded onto cognitive technology: normally executed for us unconsciously inside our brain (we are aware of our search terms, but not the retrieval process). We should be able to think much the same way if the unconconscious processing is done for us outside our brain, if we get the results fast enough.

Offloading Real-Time Cognition: Language. The evolution of language itself was probably the first way we offloaded cognitive processes that we had formerly needed to execute inside our heads: Looking for a target word? Ask someone. Of course, if your search terms are not enough to win this game of “Jeopardy” in your own head, they may not work in the head of your interlocutor either. You can write to someone else, or try looking it up in books (reading/writing being the second form of offloaded cognition), but that all takes time; and consulting a printed index for the conjunction of search terms takes still longer. Cognition is a real-time process, occurring at the speed of thought (Harnad 2003). So if mental search is to be offloaded onto external technology, the transaction needs to take place in real-time at about the speed of thought in order to feel like cognition. Language is interactive, with the dialogue occurring at about the speed of live oral conversation. That is the tempo to which both thought itself and the speed of thought have been biologically adapted. Longer turn-around times are unbiological, and certainly cannot be mistaken for online cognition. That is why googling the web feels as natural as googling your brain. Might googling the web externally reveal something about how we google our brains internally in the way that the structure of external dictionaries has cast light on our internal lexicon?

Latent Semantic Analysis. An important feature of word-associations is co-occurrence frequency: how often words occur together in text. Latent Semantic Analysis (LSA) provides a measure of this form of word-to-word “semantic” distance. We used the words in an LSA database (Maki et al. 2004) to generate the three closest LSA associates for each of a set of close to 863 words chosen from the MRC Psycholinguistic database (Wilson 1988) because they had been quantified on the three variables that had proved significant in our dictionary grounding analyses: age of acquisition, concreteness, and written frequency. Our selection criterion was that our words had to (1) have a value on each of the three MRC variables, (2) have a Wikipedia page that was devoted to them, as target word, and (3) have three closest associates in the LSA database. Starting from the smallest database (MRC) (1), and then excluding those words that failed to meet our criteria for (2) or (3), yielded the 863 target MRC words, *T*, with Wikipedia entries, plus their 3 closest LSA semantic associates, *S1*, *S2*, *S3* (2589 words). As it happened, the MRC restriction was uninformative, because none of the correlations between the success of our conjunctive and disjunctive search outcomes on the one hand, and age of acquisition, concreteness or frequency on the other hand proved significant, but the criteria did generate a small enough sample, based on well-studied words, to be tractable for this first analysis of conjunctive and disjunctive search on semantic associates, *S*, with and without the target word, *T*.

Conjunctive Versus Disjunctive Search. For each of the 863 target words, *T*, we generated a google search consisting of either the target word alone, *T*, or a disjunctive search consisting of 1, 2, or 3 of its LSA associates, *S*, as follows: (*S1*), (*S1 OR S2*), (*S1 OR S2 OR S3*). We also did disjunctive searches adding the target word: (*S1 OR T*), (*S1 OR S2 OR T*), (*S1 OR S2 OR S3 OR T*). We did the same with conjunctive searches: (*S1 AND S2*), (*S1 AND S2 AND S3*), (*S1 AND T*), (*S1 AND S2 AND T*), (*S1 AND S2*

AND S3 AND T). The results are shown in the Figure 1: As we would expect, google search on the target word, T, alone, yielded by far the best search results in terms of the number of hits of the target word's Wikipedia entry among the first ten google ranks. Also unsurprising, the hit rate is considerably lower if the LSA associates are added to the target word in the search, the more terms added the worse, and, worse when they are conjoined (AND) than when they are disjoined (OR). But what was surprising was the fact that both conjunctive and disjunctive search with 1, 2, or 3 of the target term's closest LSA semantic associates alone (i.e., the words that co-occur most frequently with the target word text), without the target word itself, generated almost no hits at all.

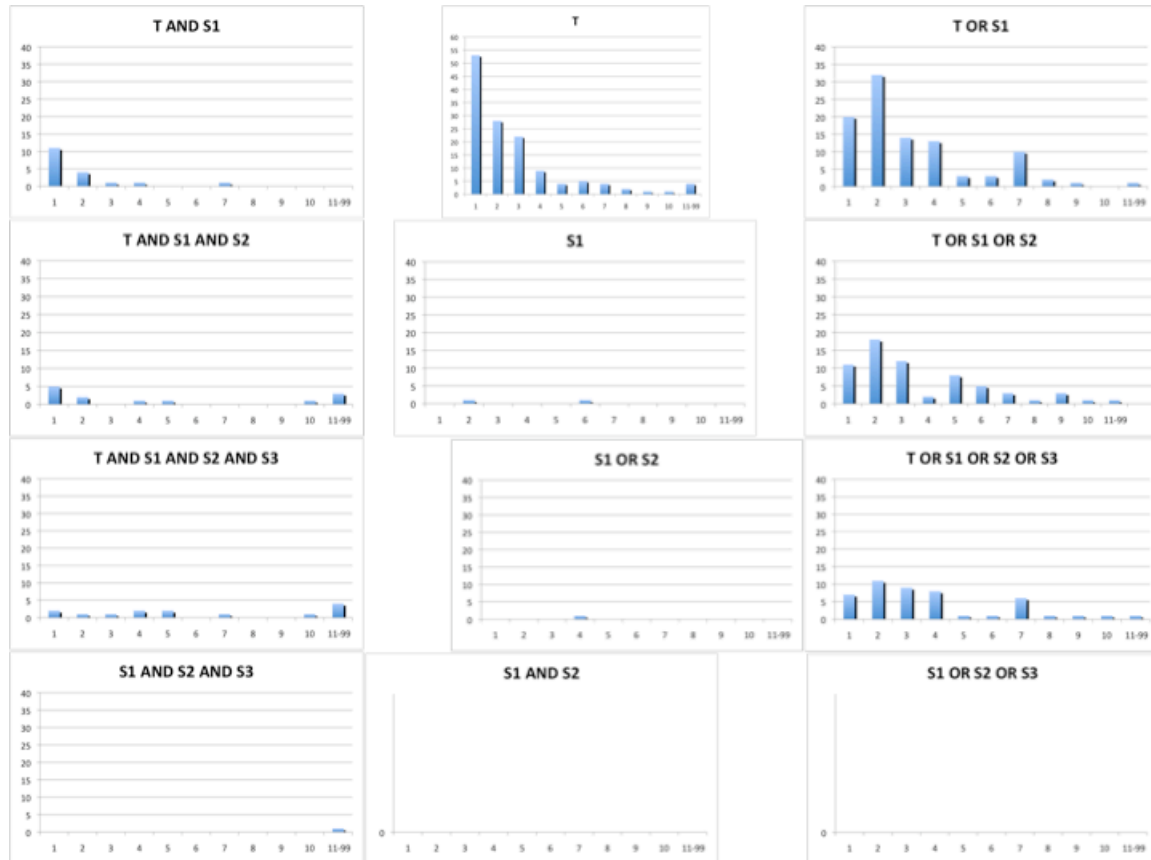


Figure 1: Success rate for conjunctive (AND) and disjunctive (OR) google searches for 863 Target words (T) each having a Wikipedia entry of its own. Success is defined by how often the Wikipedia entry for T was within the first 10 google hits. Search was either on the target word, T, alone, or on each T's three, (or fewer) closest semantic associates, S1, S2 and S3 (according to Latent Semantic Analysis [LSA] of WordNet frequency of co-occurrence with T). The high-co-occurrence words were clearly ineffective in the absence of T. Including T improves conjunctive search slightly, and disjunctive search somewhat more, but neither comes close to the effectiveness of T alone. It is unclear whether this is because (1) the LSA co-occurrence frequencies are based on a database (WordNet) that is so much smaller than google space, (2) because the google PageRank algorithm and fine-tuning heuristics are over-riding boolean disjunction and conjunction, or (3) because co-occurrence frequency plays a smaller role in the success of google searches than it does in human cognition (where it is predictive of word associations and memory retrieval), at least insofar as entries where their Wikipedia page is the criterion for success.

Local Co-Occurrence Frequency Versus Global Conjunction in Google Space. It is not clear to what extent this outcome pattern applies to conjunctive and disjunctive search in general, or it is influenced by specific heuristics used in overall google ranking (or for Wikipedia pages in particular). LSA scores are based on co-occurrence frequency, and co-occurrence means *conjunction*. A boolean conjunctive search is likewise based on conjunction. The LSA scores were based on frequency of conjunction with the target word (in the WordNet database, Fellbaum 2005). So one would have

expected the LSA conjuncts alone to do better than they did, in converging on the target word and its Wikipedia entry. It is possible that the outcome was because the WordNet database is necessarily much smaller than all of google space, and hence conjunctive google search is more representative of global conjunction effects in free text than the local WordNet database is. (We will be extending the analysis to co-occurrence analyses based on a larger chunk of google space.) It is also possible that in google in general (or for google Wikipedia entries for direct target terms), there is a tension between boolean conjunction and the PageRank algorithm or some of its heuristic fine-tuning. It is known that LSA scores are correlated with and predictive of human word-association performance and other psycholinguistic measures. So this divergence is worth investigating further, and we plan to do so with other forms of semantic relations, in addition to LSA co-occurrence, as derived from WordNet (Fellbaum 2005), such as synonymy, hyperonymy, and mereonymy.

Semantic Web Tagging. One widely held hope is that google space in general, and digital free text in particular, will be made more intelligently navigable if documents, passages in documents, phrases and words are explicitly marked up by creators and users with their semantic descriptors. But what are semantic descriptors? Like dictionary definitions, they too are just words, combined and recombined. Although the present pilot study failed to detect any significant correlations with the variables that characterized the grounding vocabulary in our dictionary studies (age of acquisition, concreteness and frequency), our sample was small, and we will extend it to a much larger one, using the British National Corpus (Aston & Burnard 1998). The hypothesis is still that there is a small core of words, implicitly encoded in sensorimotor representations based on direct experience. These ground the meanings of most other words, which we instead learn through explicit verbal description and definition consisting of recombinations of already grounded words, so that their meaning is explicitly represented verbally. Those same grounded words are also likely to play a special role in boolean search of google space. Our hunch is that both combinatory verbal representations of word meanings and boolean associative search across verbal databases can be readily offloaded onto cognitive technology, but the encoding of the grounding terms cannot.

REFERENCES

- Aston, Guy & Burnard, Lou (1998) *The BNC handbook: exploring the British National Corpus with SARA*. Edinburgh University Press. Database: <http://www.natcorp.ox.ac.uk/>
- Blondin Masse, Alexandre et Chicoisne, Guillaume et Gargouri, Yassine et Harnad, Stevan et Picard, Olivier et Marcotte, Odile (2008). [How Is Meaning Grounded in Dictionary Definitions?](#) In TextGraphs-3 Workshop - 22nd International Conference on Computational Linguistics.
- Dror, I. and Harnad, S. (2009) [Offloading Cognition onto Cognitive Technology](#). In Dror, I. and Harnad, S. (Eds): *Cognition Distributed: How Cognitive Technology Extends Our Minds*. Amsterdam: John Benjamins
- Fellbaum, C. (2005) Theories of human semantic representation of the mental lexicon. In: Cruse, D. A. (Ed.), *Handbook of Linguistics and Communication Science*, Berlin, Germany: Walter de Gruyter, 1749-1758.
- Harnad, S. (2003) [Back to the Oral Tradition Through Skywriting at the Speed of Thought](#). Interdisciplines. Presses de l'enssib.
- Harnad, S; Blondin-Massé, A, St-Louis, B; Chicoisne, G; Gargouri, Y; & Picard, O. (2008) [Symbol Grounding, Turing Testing and Robot Talking](#). RoadMap Workshop on Action and Language Integration, Rome on 22-24 September 2008.
- Maki, William S., McKinley, Lauren N & Thompson, Amber G. (2004) [Semantic distance norms computed from an electronic dictionary \(WordNet\)](#) *Behavior Research Methods, Instruments, & Computers* 36 (3): 421-431 (Database: <ftp://ftp.ttu.edu/pub/maki>)

Wilson, M.D. (1988) The MRC Psycholinguistic Database: Machine Readable Dictionary. *Behavioral Research Methods, Instruments and Computers*, 20(1), 6-11.