Local Associations and Global Reason: Fodor¹s Frame Problem and Second-Order Search² COGNITIVE SCIENCE QUARTERLY 2:2:2002 :115-140

Global Abductive Inference and Authoritative Sources, or, How Search Engines Can Save Cognitive Science*

Andy Clark

Philosophy/Neuroscience/Psychology Program
Department of Philosophy
Washington University
St. Louis, MO 63130
e-mail: andy@twinearth.wustl.edu

* Thanks to Luis Rocha and the information technology group at Los Alamos National Laboratory for useful materials and conversations. Thanks also to Bill Wimsatt, Bill Bechtel, Jesse Prinz, Dominic Murphy and the participants in the History and Philosophy of Science Seminar at Washington University in St. Louis.

Abstract

Kleinberg (1999) describes a novel procedure for efficient search in a dense hyper-linked environment, such as the world wide web. The procedure exploits information implicit in the links between pages so as to identify patterns of connectivity indicative of “authorative sources”. At a more general level, the trick is to use this second-order link-structure information to rapidly and cheaply identify the knowledge-structures most likely to be relevant given a specific input. I shall argue that Kleinberg’s procedure is suggestive of a new, viable, and neuroscientifically plausible solution to at least (one incarnation of) the so-called “Frame Problem” in cognitive science viz the problem of explaining global abductive inference. More accurately, I shall argue that Kleinberg’s procedure suggests a new variety of “fast and frugal heuristic” (Gigerenzer and Todd (1999)) capable of pressing maximum utility from the vast bodies of information and associations commanded by the biological brain. The paper thus takes up the challenge laid down by Fodor ((1983)(Ms)). Fodor depicts the problem of global knowledge-based reason as the point source of many paradigmatic failings of contemporary computational theories of mind. These failings, Fodor goes on to argue, cannot be remedied by any simple appeal to alternative (e.g. connectionist) modes of encoding and processing. I shall show, however, that connectionist models can provide for one neurologically plausible incarnation of Kleinberg’s procedure. The paper ends by noting that current commercial applications increasingly confront the kinds of challenge (such as managing complexity and making efficient use of vast data-bases) initially posed to biological thought and reason.

Introduction: The Frame Problem (maybe).

Back in 1969, McCarthy and Hayes coined the term "the frame problem." Since then the term has meant many things to many people, none of them good (for reviews, see Pylyshyn (1987), Ford and Hayes (1991), Ford and Pylyshyn (1996)). What it meant to McCarthy and Hayes (1969), had to do largely with the updating of a system's understanding of the world following activity by the system. At the heart of this problem seems to lie a more general, and painfully intractable, puzzle. It is the puzzle of finding the *right* stuff (information, data) to consider (update, or use in reasoning) at the *right* time. The puzzle is intractable because in many cases there is no obvious sub-set of data or information to which we can reasonably restrict the search. What is required, it often seems, is to be sensitive to the entire contents of the data-base, so as to make an inference to the best explanation, to choose an action, or to fix a belief, in a way that is maximally consistent with the global data-base. This is the version of the frame problem that has Fodor "worried half to death" ((Ms) p. 37)), and that formed the pessimistic centerpiece of Fodor's "first law of the non-existence of cognitive science" viz "the more global … a cognitive process is, the less anybody understands it" (Fodor (1983) p. 107).

To be clear, then, it is this problem – the problem of global abductive inference, as I'll later cast it, that I shall be targeting in the text. Whether this is indeed the frame problem as understood by McCarthy and Hayes doesn't really matter. At the very least, it is a problem that arises in the vicinity, and one that Fodor sees as potentially fatal to the ambitions of current cognitive science.

The discussion goes like this. In section 1 (next) I isolate the particular aspect of the particular problem I want to target viz (roughly) the sub-problem of how to find the most relevant body of beliefs and information in a massive knowledge-base. This problem, I shall show, constitutes a very large piece of Fodor's puzzle. In section 2, I display a potent and ingenious solution developed by Jon Kleinberg as a means of effective search in a dense hyper-linked-environment like the world wide web. Section 3 returns to cognitive science territory and shows first, that the simple appeal to connectionist modes of processing and storage does not solve so much as relocate Fodor's problem, but second, that it relocates it in a way that suggests a neurally plausible incarnation of Kleinberg's strategy. I end (section 4) by briefly commenting on the wider significance of this result: that successful negotiation of the world wide web now raises many of the same problems that once confronted technology is thus increasingly well-positioned to recapitulate phylogeny.

1. Commonsense, Modularity, and Global Abductive Inference.

Cognitive science, according to Fodor, has had one – and only one – decent idea about how the mind works. It is the idea, roughly, that

> What makes minds rational is their ability to perform computations on thoughts, where thoughts, like sentences, are assumed to be syntactically structured, and where "computation" means formal operations in the manner of Turing.
>
> Fodor (1998) p. 205.

Thus it is possible to understand how a merely physical system might be set up so as to abide by the rules of deductive inference: to token (to use Fodor's own (op cit) example) P and Q, if and only if P is true and Q is true, to infer P from P & Q and so on. This, Fodor frequently reminds us, is a quite considerable advance; it is the heart of the classical computational theory of mind which (still according to Fodor) offers our only scientific foothold into the realms of normativity and rationality. But it is a foothold that cannot, Fodor believes, get us to the top of the mountain. In fact, it leaves us, infamously, stranded in the foothills of input systems and deductive inference, with the peaks of central processing, belief fixation, and global abductive inference (more on which shortly) towering tauntingly above us.

The trouble – *all* the trouble, it seems, with the exception of conscious awareness about which we shall remain silent – has to do with non-locality. It is where processing, inference and recall need to be sensitive to properties not of single tokens, but of entire bodies of knowledge ("whole belief systems" in Fodor's (op cit, p. 206) phase) that Turing style computation threatens to fall short. There exist several sub-problems hereabouts, all of which turn, in somewhat different ways, on the need to be sensitive to non-local properties of neural representations – see e.g. Fodor (Ms) Ch. 2. In this paper, I am concerned with just one (extremely central) such sub-problem: the challenge of tractable search and recall given an extremely large data-base. It is this problem that is most intimately tied up with the issues concerning belief fixation and global abductive reasoning. And it is this problem that Fodor typically glosses – for better or for worse – as the "frame problem in cognitive science." In Fodor's own words:

> "The frame problem is a name for one aspect of the question how to reconcile a local notion of mental computation with the apparent holism of rational inference; in particular will the fact that information that is

relevant to the optimal solution of a abductive problem can, in principle, come from anywhere in the network of one's prior epistemic commitments."

Fodor (Ms) p. 41.

Here is one quick way to appreciate the worry. Consider a system that knows a great deal (like us), and that receives a piece of apparent information from the world. Let's imagine, to be concrete, that it's visual system seems to be telling it that there is a pink elephant hovering over the mantelpiece. What should it believe? The point, frequently stressed by Fodor (e.g. (1983) p. 102), is that in deciding what to believe the system may draw on beliefs and information of many kinds, and that it need not – and should not – simply take the visual percept at face value. For example, if you know that you tend to see pink elephants every new millennium (ok, year), and that the effect is closely scaled to your champagne consumption, you might well come to believe not that you are indeed seeing a pink elephant, but that you should cease partying.

Now this is *not* a problem, notice, for the input system that (in Fodor's view) computed the visual image of the pink elephant in the first place. Such systems correspond, in Fodor's cognitive architectural ontology, to *modules*, which can get by using good, old-fashioned (local, syntactic, Turing-friendly) methods of computation. Such modules can even, at times, compute results which depend on the global properties of the restricted data-base (for example, the "pink elephant" perception may have been the *globally simplest* interpretation of the whole set of data available to the visual processing system). What matters is that the input systems can afford to rely on local, syntax (form) based operations. This is fine in a module, which by definition (see Fodor (1983) p.) has an encapsulated and hence realistically navigable data-base. The problem of *what* data to consider never arises since it is not fatal to assume that *all* the data, in the restricted module, is potentially relevant and to proceed accordingly. Cognitive modules, with their encapsulated data-bases, thus do not confront the frame problem (as Fodor understands it) since:

> "Frame problems and relevance problems are about how deeply, in the course of cognitive processing, a mind should examine its background of epistemic commitments. Modular problem solving doesn't have to worry about that sort of thing because its searches are constrained architecturally; what is in the data-base can count … and nothing else counts as relevant."
>
> Fodor (Ms) p. 61

The problem of global abductive inference is best appreciated against this backdrop. It is the problem, in a super-compressed nutshell, of how to do inference to the best explanation in a way that is sensitive to whatever is *most relevant* in the *massive* body of belief and knowledge that underlies commonsense thinking and reasoning. Abductive reasoning, is, roughly, reasoning in which you explain, or understand, an event (an apparent event, in the pink elephant case) by coming up with the set of antecedent conditions which – given what you already know and believe – best account for the event. For example, you find the 'phone off the hook and infer that the neighbor's cat has once again invaded your apartment. This conclusion is obviously not deductively implied by the evidence and your background beliefs. But given the evidence and your background beliefs, it is the best explanation of the ’phone’s being off the hook. Such modes of reasoning, as Peirce and others have persuasively argued, characterize both commonsense and scientific thought and reason. Classical computational approaches confronted a major challenge, hereabouts. Ordinary commonsense reasoning, as Dreyfus and others were quick to notice, seemed (and still seems) largely intractable given Turing style computational strategies. The problem was two-fold: a problem of search, and a problem of “weighing.” The problem of search was simply how to find, given a massive (indeed, seemingly boundless) set of explicit and implicit beliefs, the sub-set most relevant to the problem at hand. The problem of ‘weighing’ was the problem, even given such a sub-set, of deciding which beliefs and items of knowledge to rely on the most. Peter Lipton (1991) thus divides (this version of) the frame problem into two subcomponents which he calls “epistemic filters.” One filter negotiates the massive background data-base to generate candidate abductive explanations. The other selects the best explanation from this set. I’ll concentrate largely on the first filter, though the solution to this problem (as we’ll see) actually goes a long way towards solving the other problem as well.

The goal, then, is to directly address the worry that central cognition (defined as those processes especially implicated in commonsense and scientific thought and reason) relies heavily on abductive inference and that “reliable abduction may require, in the limit, that the whole background of epistemic commitments be somehow brought to bear in planning and belief fixation” (Fodor, (Ms) p. 37). What gives this worry its practical bite is that, due to the local, syntactically-driven nature of Turing-style computation, “*feasible* abduction requires…that not more than a small subset of even the relevant background beliefs [be] actually consulted”. The frame problem, for Fodor, just *is* the problem of “how to make abductive inferences that are both feasible and reliable” (quotes from Fodor (Ms) p. 37, emphasis mine).

This problem is, Fodor suspects, just about terminally embarrassing for familiar cognitive scientific accounts of central processing. It explains, Fodor believes, our persistent inability to build a halfway decent household robot. This is no surprise if, as Fodor claims, “we don’t have a theory of commonsense reasoning that would survive scrutiny by an intelligent five year old” (Fodor (1998) p. 206)). Some have thought – and not without reason – that Fodor is right about classical computational approaches but that connectionism (artificial neural networks) either sidesteps or solves the frame problem. But for once I (almost) agree with Fodor. Connectionism doesn’t solve the problem, though it does relocate it in a potentially fruitful way. In fact, it relocates it – as I’ll argue later on (section 3) – in exactly the right place for a neurologically plausible incarnation of the strategy I am about to describe. First, though, let’s meet the strategy on its somewhat surprising home-ground: the problem of effective search in a dense hyperlinked environment. Time, then, to consider the more familiar problem of navigating the world-wide web.

2. Link Structures: Using Implicit Knowledge to Tame the Web.

The development of web search engines, along with the co-evolutionary explosion of web sites, has led to a problem with which we are all depressingly familiar. The information we need – the precise, exact thing we most need to know right now – is probably out there somewhere, packaged and waiting, *but we just can’t find it*. Granted, there are many search engines: but they seldom provide the fast, appropriate information retrieval we need. Indeed, they often retrieve (even when used properly) voluminous junk, and have a regrettable tendency to miss the good stuff altogether. There is a sense in which this is not surprising, for the problem they confront is formidable. There are often literally millions of pages whose contents look superficially relevant – especially given that the usual test for relevance is dumb syntactic matching: the search engine seeks pages that either contain, or are indexed as containing, tokens of the specific string or strings entered by the user. The situation is worsened by the unplanned, anarchic nature of the web itself – there is little deliberate global organization of the kind that might be useful in streamlining search.

Web-based search thus presents a familiar difficulty which Kleinberg dubs "the abundance problem";

> *The Abundance Problem: The number of pages that could reasonably be returned as relevant is far too large for a human user to digest.* To provide effective search methods under these conditions, one needs a way to filter, from among a huge collection of relevant pages, a small set of the most "authoritative" or "definitive" ones.
>
> Kleinberg (1999) emphasis in original

The abundance problem does not arise for more local, restricted kinds of search, where the number of returns from even a dumb syntax-based search remain tractable. It is a problem only with *global* searches – ones confronting the whole WWW – since:

> Global approaches involve basic problem of representing and filtering large volumes of information… This is in contrast to *local* approaches [where] the amount of data is much smaller and…a different set of considerations dominates.

Certain parallels between the problems outlined in Section 1 and the abundance problem for web-based search are already clear. In each case there is a formidable body of potentially relevant data, yet we want a search procedure that retrieves only some especially relevant subset. And in each case the problem is exacerbated by the dumb, purely local and syntactic nature of the mechanistic search procedures available. Finally, there is a clear contrast, in each case, between procedures that, though dumb, local and syntax-driven, are *good enough* for searching restricted data-bases (encapsulated modules, sets of web pages belonging to a single site or intranet) and the kinds of procedures needed to tackle the more global case.

The good news is that the abundance problem, at least, has a solution: one currently the object of much ongoing research within the hypertext information retrieval community. The key trick, it seems, is to focus attention not on the *content* of the pages so much as on the *structure* of *links* between pages. The hyperlink structure itself – the way different pages link to and from each other – turns out to be a treasure house of implicit knowledge concerning which pages are most central and authoritative regarding a given topic. Kleinberg's achievement is to devise a set of algorithms that extract and utilize the knowledge thus implicit in the web of connectivity. Formally, the algorithms construct a focused subgraph of the WWW and then exploit a number of additional tricks and heuristics to further sort and refine the results. Interestingly, the actual *contents* of specific pages are consulted only in the very first stages of the complex procedure. Most of the real work, as noted, depends purely on the analysis of the patterns of interconnectivity (link structure) between pages. The search is thus entirely dumb and 'syntax-driven' (in Fodor's sense of syntax) but the relevant syntactic features are largely *higher-order*: they concern the shape not of the knowledge-structures themselves but of the hyperlinks between them. My suggestion – to be laid out in Section 3 following – will be that this trick (of looking at link structures rather than knowledge structures) is also exploited by biological brains and that it is what enables them to solve the formidable problems of search involved in global abductive inferences. But first, let's look at Kleinberg's solution in a little more detail.

Kleinberg begins by dividing the types of query a user might have into three broad types: *specific* queries, such as "Does Netscape support the JDK 1.1. code signing API", *broad topic* queries, such as "find information about Java", and *similar page* queries, such as "find pages similar to java.sun.com". Specific queries are set aside, as they do not raise the abundance problem with which Kleinberg is concerned. Broad topic queries are thus the main focus, but the same general procedure (of attending to hyperlink structure rather than primary content) turns out to have application to the similarity problem. (Here again, there is non-accidental overlap with Fodor's concerns: reasoning by analogy, Fodor (1983, p. 107) suggests, depends upon just the kind of process (intelligent search in a massive space) that is characteristic of both science and central systems, and about which cognitive science remains firmly in the dark).

To illustrate the abundance problem, Kleinberg takes the simple query string "Harvard". It so happens that:

> there are over a million pages on the WWW that use the term "Harvard" and www.harvard.edu is not the one that uses the term most often, or most prominently, or in any way that would favor it under a text-based

> ranking function. Indeed, one suspects there is no purely *endogenous* measure of the page that would allow one to properly assess its authority.
>
> Kleinberg (1999) p. 2

Other examples: to search for “search engines” – this is especially tough, since many of the most authoritative pages (Yahoo, Excite, Alta Vista) do not use the term on their home pages; to search for very broad topics, such as “censorship” – this tends to return a hodge-podge of largely non-authoritative sources, (ditto for searches such as “Gates”, seeking authoritative information about Bill Gates and Microsoft – again, the pages retrieved tend to be non-authoritative). Standard searches thus tend to be both inefficient (return too much) and insufficiently intelligent (despite returning too much, they often miss – or return way down in the list – the most relevant and authoritative sites).

Kleinberg’s procedure *starts*, nonetheless, with a dumb standard search. It first collects – for some tractable number ò – the top ò highest ranked pages returned (for some broad query) by a standard search engine. This is the only time text-based search is invoked, and it delivers a “root set R” of pages – a set which, we just argued is quite likely to *fail* to contain the pages in which we are in fact most interested. The next step is to seek a set of pages that is, in fact, more likely to contain the pages we need. The key assumption here is that the authoritative pages, though perhaps missing from the root set (the top ò pages returned by standard search) is probably at least *pointed to* by one or more pages in R. R. is thus expanded to include any page pointed at by a page in R, and any page that points to a page in R (with some key restrictions to keep things manageable[1]).

Now things get interesting. What you have in hand is a directed graph G, in which nodes correspond to pages and directed edges reflect links between specific pages. Call the number of nodes a node is linked to the "out-degree" of the node and the number of nodes that link to a node the "in-degree" of the node. The question is how, by exploiting the link-structure information given by this directed graph, to discern the authoritative sources among the set of nodes (pages). An obvious move is to consider the relative in-degrees of the nodes; a procedure which works as long as the most authoritative sources are the most widely indexed (have the greatest number of links – subject to the prunings mentioned in footnote 1 – *to* them).

The trouble here is that mere popularity does not authority make. Indeed, some sites are almost universally popular, and have enormous in-degrees regardless of topic: amazon.com is a prominent example. Further filters are clearly required. Kleinberg notes that if a site is indeed authoritative with respect to a query, we may expect not *just* that many other sites in G link to it, but that there be certain pages – "hub pages" – that have links to multiple relevant authoritative pages. Good authorities, likewise, will be pointed to by multiple such "hub pages". The heart of Kleinberg's procedure is an algorithm that computes, from the link structure of G, this mutually inter-defined set of hubs and authorities. The algorithm works by computing the "eigenvectors of certain matrices associated with the link graph" (op cit, p. 1). To fully display the algorithm requires the introduction of several notions from linear algebra, and I here refer the reader to the comprehensive account in Section 3 of Kleinberg's paper. Let us pause, though, to appreciate the results. Here, for example, are the top search results for the query strings "java", "censorship", "search engines" and "Gates" (the numbers on the left indicate the overall strength of the "authoritativeness" rating):

(java) Authorities

.328 http://www.gamelan.com/	Gamelan
.251 http:// java.sun.com/	JavaSoft Home Page
.190 http:// www.digitalfocus.com/digitalfocus/faq/howdoi/html	The Java Developer:How Do I…
.190 http:// lightyear.ncsa.uiuc.edu/~srp/java/javabooks.html	The Java Book Pages
.183 http://sunsite.unc.edu/javafaq/javafaq.html	comp.lang.java FAQ

(censorship) Authorities	
.378 http://www.eff.org/	EFFweb – the Electronic Frontier Foundation
.344 http://www.eff.org/blueribbon.html	The Blue Ribbon Campaign for Online Free Speech
.238 http://www.cdt.org/	The Center for Democracy and Technology
.235 http://www.vtw.org/	Voters Telecommunications Watch
.218 http:// www.aclu.org/	ACLU: American Civil Liberties Union

(“search engines”) Authorities	
.346 http:// www.yahoo.com/	Yahoo!
.291 http:// www.excite.com/	Excite
.239 http:// www.mckinley.com/	Welcome to Magellan!
.231 http:// www.lycos.com/	Lycos Home Page
.231 http:// www.altavista.digital.com/	Alta Vista: Main Page

(Gates) Authorities	
.643 http:// www.roadahead.com/	Bill Gates: The Road Ahead
.458 http:// www.microsoft.com/	Welcome to Microsoft
.440 http:// www.microsoft.com/corpinfo/bill-g.htm	

Kleinberg (1999) p. 11-12

To compare this to text-based search, it is useful to note that almost *none* of these pages appeared in the root set R (the top ò pages returned by text-based search). They instead first appeared in the subgraph G obtained by expanding the root set along the links that enter and leave it, and obtained prominence only by the further computation of likely hubs and authorities carried out using the fine-grained link structure information. In fact, the only single page in the above list that was returned by the original text-based search (as a member of R) was www.roadahead.com, returned as the 123rd choice of Alta Vista!

Kleinberg stresses that almost all of this remarkably effective procedure operates by ignoring actual page content, focusing instead on the link structure alone. It is this trick, of looking at the information-about-information (second-order information) implicit in the link structures that may help solve some aspects of the more general problem of global abductive inference. What the shift to looking at secnd order information really does, I suggest, is to create a useful, low-dimensional reflection of the high-dimensional knowledge-space. Kleinberg sums it up nicely:

> our algorithm produces pages that can legitimately be considered authoritative with respect to the WWW as a whole, despite the fact that it operates without direct access to any large-scale index of the WWW. Rather, its only “global” access to the WWW is through a text-based search engine such as Alta Vista, from which it is very difficult to directly obtain reasonable candidates for authoritative pages on most queries. What the results imply is that it is possible to reliably estimate certain types of global information about the WWW using only a standard search engine interface; a global analysis of the full WWW link structure can be replaced by a much more local method of analysis on a small focused subgraph.
>
> Kleinberg (1999) p. 12

Kleinberg’s results are clearly impressive, and make significant headway with a very real and practical problem of global search. But can they truly help us to understand how biological brains might be solving some of the puzzles highlighted in Section 1: puzzles which lie, we saw, somewhere in the vicinity of the infamous frame problem itself? The answer, I’ll next argue, is a tentative but tantalizing ‘yes’.

3. Neural Networks, Gating, and the Frame Problem Re-visited.

Three hurdles stand in the way of any simple appeal to Kleinberg's procedure (henceforth, KP) to solve the kinds of problems about human cognition highlighted in Section 1. They are:

(1) The fact that KP is defined over familiar symbolic encodings of information (web pages), whereas biological information storage plausibly involves quite different modes of encoding.

(2) The relative balance, in the problems targeted in Section 1, between issues concerning search and retrieval and issues concerning the evaluation of the information retrieved (Lipton's "second filter").

and

(3) The role of human intelligence and intuition in *setting up* the link structures on which KP depends (perhaps it is only because human brains can *already* solve the frame problem that we can set up the link structures that allow KP to work, thus rendering KP circular as a solution to the frame problem).

I shall address each worry in turn. First, then, the question of the knowledge structures over which KP operates. These are familiar (indeed, paradigmatic) symbolic encodings. Yet the biological brain, many of us believe, eschews such symbol-mongering in favor of various kinds of distributed, neural population based encoding. Some have argued, in fact, that once you move to this kind of encoding (and processing) system, the frame problem as Fodor imagines it simply does not arise. If this is true, KP is cognitively spurious. If it is false (it *is* false), we will need to see whether KP can also get a grip in a more "connectionist" kind of setting.

Does a broadly connectionist approach to encoding and processing already solve (Fodor's version of) the frame problem? That version, recall, was at heart a problem about non-demonstrative inference and the potency of abductive reason in the context of a large, non-modular knowledge-base. More precisely, that *part* of the problem with which we are concerned is what Haselager (1997) calls the descriptive and computational part viz "how *part* of what one knows influences what one *ends up* believing" (op cit, p. 105). Paul Churchland, however, has suggested that this aspect of the frame problem, at least, is nicely taken care of by the adoption of an alternative ("connectionist") mode of encoding and retrieval. In Churchland's view:

> "the depiction of one's knowledge as an immense set of individually stored "sentences" raises a severe problem concerning…relevant retrieval…How is it one is able to retrieve, from the millions of sentences stored, exactly the handful that is relevant to one's current predictive or explanatory problem, and how is it one is generally able to do this in a few tenths of a second? This is known as the "frame problem" in AI and it arises because…a long list of sentences is an appallingly inefficient way to store information."
>
> P.M. Churchland (1989) p. 155-56

Connectionist networks – and, let us assume for the sake of argument[2], the brain – store information in a radically different way. A full, or even adequate, account of all this would be out of place here (see e.g. Clark (1989)(1993) or Churchland (1989)(1995)). But the key features that are supposed to help solve the frame problem are easily explained. Connectionist networks replace serial search with spreading activation in a massively parallel system. Within that system, knowledge is stored not as discrete symbolic strings but as weights between unstructured nodes. The weights (numerical quantities, positive and negative, which differentially affect the flow and strength of node activations) allow the system to *retrieve* information in the form of activity-vectors across a whole population of nodes. Such retrieval is typically fast and nicely tuned to the impinging stimulus (retrieval cue). Thus suppose the retrieval cue is a visual input as of a kitchen full of smoke. Spreading activation through a massively parallel network is what allows us, according to Churchland, to understand "at a glance why…the kitchen is filled with smoke: the toast is burning!" (Churchland (1989) p. 199). It is worth dwelling on the proposal. The idea is that the frame problem, conceived as a problem of "speed-of-relevant-access" (op cit, p. 178) is directly solved by the connectionist approach. Here is the passage in full:

> "A network the size of a human brain – with 10^{11} neurons, 10^{3} connection on each, 10^{14} total connections, and at least 10 distinct layers of hidden units – can be expected, in the course of growing up, to partition its internal vector spaces into many billions of functionally relevant subdivisions, each responsive to a broad but proprietary range of highly complex stimuli. When the network receives a stimulus that falls into one of these classes, the network produces the appropriate activation vector in a matter of only tens or hundreds of milliseconds, because that is all the time it take for the parallel-coded stimulus to make its way through only two or three or ten layers of the massively parallel network to the functionally relevant layer that drives the appropriate behavioral response. Since information is stored not in a long list that must somehow be searched, but rather in the myriad connection weights that configure the

> network, relevant aspects of the creature's total information are automatically accessed by the coded stimuli themselves.

The problem, I think, should be clear. In order to press a solution to the frame problem so *directly* from connectionist encodings, Churchland has had to innocently rig the game. He has assumed that first, the brain is already divided into "billions of functionally relevant subdivisions" and second, that each such subdivision gets to "see" the current stimulus so as to have the chance to (as it were) resonate to it or not. But neither assumption, it seems to me, is realistic. The choice of *billions* of relevant subdivisions is not accidental. Churchland is not here imagining a few thousand neuro-anatomically distinct areas (a more realistic bet): he is imagining enough *functionally* distinct subdivisions (not necessarily neuro-anatomically distinct) to devote a single subdivision to each distinct and integrated body of knowledge we possess. In that respect, the solution exploits something not unlike good old fashioned frames (Minsky (1975)) or scripts (Schank and Abelson (1975)), but with the additional resource of connectionist (parallel, distributed) encoding and retrieval. But there are familiar problems here. Either there are uncountably many such subdivisions, with massive re-duplication of information, or the ones that exist are densely cross-referenced. It is uncontentious, I believe, that the latter is the only practical possibility. But as Fodor is not slow to notice, there is in that case:

> No reason to doubt that…the system of cross-referencing would imply a graph in which there is a route…from each point to any other. But now we have the frame problem all over again: Which such paths should actually be traversed in a given case of problem-solving…? All that has happened is that, instead of thinking of the frame problem as an issue in

> the logic of confirmation, we are now invited to think of it as an issue in the theory of executive control…
>
> Fodor (1983) p. 117

It may seem as if the connectionist solution sidesteps this problem by allowing *all* the brain's knowledge-level resources a chance to be "automatically accessed by the coded stimuli" (see earlier quote). But this is not at all the case. Instead, incoming information is rapidly routed, filtered and transformed as it flows around the brain. The incoming stimulus is simply not encountered by most of our neural resources. Instead, information flows in a way which is well-suited (typically) to the problem at hand. What Fodor calls the problem of "executive control" – I'll just call it the "routing problem" – is thus where the frame problem, in connectionist guise, spills out from under the rug. Get the flow of information-retrieving information (stimulus cues, suitably transformed) right, and connectionist modes of encoding and storage do indeed result in gains of speed and efficiency. But what modulates the flow? Something more is needed. Something simple connectionist models do not – pace Churchland's optimistic prognosis – provide (for a similar diagnosis, see Haselager (1997, Chapter 5) and Fodor (Ms, Chapter 3).

What is needed, I suggest, is some combination of (i) a certain kind of second-generation connectionist architecture (see below) and (ii) the exploitation, within such an architecture, of the kind of second-order (link-structure) information highlighted in KP. Here's what I have in mind.

Van Essen et al (1994) is a neuroscientifically plausible, connectionist-style account that describes "an explicit mechanism for dynamically regulating the flow of information within and between cortical areas" (op cit, p. 271). Such a mechanism is required, we are told, due to:

> Computational considerations relating to (1) the vast amounts of data continuously impinging on the nervous system, (2) the finite computational resources that can be dedicated to any given task and (3) the need for highly flexible linkages between a large number of physically separate modules.
>
> Van Essen et al (1994) p. 271

What Van Essen and his colleagues offer, in fact, is a working, neurobiologically plausible model which accounts for the way appropriate information is channeled to the *right* functional subdivisions of cortex at the right time. They are thus proposing a solution to the problem which – I suggested – afflicts Churchland's appeal to connectionism as a means of dissolving the frame problem. At the heart of Van Essen et al's picture is the idea that many neurons and neuronal populations serve not as direct encodings of knowledge or information, but as (dumb) middle managers routing and trafficking the internal flow of information between and within cortical areas. These "control neurons" serve to open and close channels of activity, and allow for the creation of a kind of instantaneous, context-sensitive modular cortical architecture: control neurons weave functional models "on the hoof", in a way sensitive to the effects of context, attention and so on. Van Essen et all develop a detailed, connectionist model of simple versions of such circuits. Related proposals include Edelman and Mountcastle's work on "reentrant processing" in which feedback and feedforward pathways are used to control and co-ordinate activity in multiples sites, and Damasio and Damasio's (1994) notion of "convergence zones", which are neuronal populations which likewise initiate and co-ordinate activity in multiple neuronal groups. It is worth noting that in none of these cases are the routing and co-ordinating circuits themselves highly intelligent problem-solving homunculi: instead, they are themselves just more, relatively dumb, connectionist-style resources. How, then, do they get the routings right? In the case of Van Essen et al's "Shifter circuits" (the circuits in which control neurons route incoming information) the required knowledge (about what inputs to send where and when to do so) is presumed to be innate. This is perhaps plausible in the case of some aspects of visual attention, recognition and motor control (the targets of their simple model). But it is less plausible for higher level cognition: a domain in which they nonetheless comment that:

> It requires only a modest conceptual leap to suppose that analogous routing strategies may be used to control the flow of information in whatever central structures are used to represent semantic information and other high-level abstractions that are the coinage of cognitive function.
>
> Van Essen et al (1994) p. 298

As things stand, however, the modest leap looks a trifle immodest: a touch flagrant, if it requires that *this* routing information be innate, and a touch mysterious (given the absence of a mechanism for learning appropriate routings) if not.

Perhaps KP can be of help. The problem, it seems, is how to learn, given a developing, experience-driven core of worldly knowledge, what aspects of that knowledge should be pulled together into an instantaneous, soft-wired module, given a specific current stimulus or input. As we saw, the stimulus itself must indeed do some of the work here: it must effectively select, without the intervention of *intelligent* middle-managers, a set of sub-populations of neurons to which to present its case. What this corresponds to, in the case of web-based search, is the need to create – on-the-hoof – a suitably restricted subgraph (G) of the web. As Fodor puts it, it is thus "*unstable instantaneous* connectivity that counts" (1983, p. 118) – connectivity that "changes from moment to moment as dictated by the interaction between the program that is being executed and the structure of the task in hand" (op cit).

KP goes some way towards showing how this might be achieved. The first step, recall, is to allow a relatively dumb process to activate an artificially limited set of knowledge structures: pages in KP, functional subdivisions of cortex in the neural case. These resources will be far, far from adequate, but may at least be presumed to have links to the resources (pages, knowledge-structures) we really need to recruit. At this point KP proper kicks in, operating not on the content of these first-pass-identified resources, but on their second-order properties: this is the link-structure information, in the web case, and the *patterns* of larger-scale associative links in the neural case (still assuming a broadly connectionist encoding). These second-order features are a tractable search object, and afford a kind of content-free sparse image of the organization of the associative links leading to and from the first-pass-identified resources (the knowledge structures in the root set R). Using this information, a neural implementation of KP could then identify a much *better* set of resources (knowledge-structures) – the neural analogues to the "hubs and authorities" (see Section 2 above) relevant given a certain input.

The first-pass search (which, in KP, used a standard text-based search engine constrained to deliver only a fixed number (ò) of returns) corresponds, I suggest to the simplest kind of connectionist retrieval mode – the kind appealed to by Churchland. It identifies some neural populations that resonate to the stimulus, but – given that the input cannot be shown to the whole brain, and given the relatively dumb, pattern-matching nature of the search – it is not expected that the resources we most need to recruit will always be discovered at this stage. Instead, something like link-structure information will often need to be used, in a second-pass search, to identify (without yet accessing any actual knowledge-structures) a better set of resources (corresponding to the subgraph G) whose link structure patterns allow the identification of hubs and authorities.

The big question, obviously, is can the biological brain actually compute the kind of link-structure (associative link-structure) information needed to drive something like KP. The immediate problem is again one of locality. If we eschew – as I think we must – the image of a kind of overseeing executive in the brain, able to obtain link-structure information by direct inspection, then we need to understand how locally-driven kinds of spreading activation might nonetheless be sensitive to information implicit in the higher level pattern of links itself. What KP provides is a nice existence proof that *if* some relevant analogue to link-structure information is available, it can be used to make truly important inroads into the search-based aspects of the frame problem and the problem of abductive inference. This is itself a highly significant result, and reason enough, I believe, to take very seriously the project of discovering a locally-computable means of identifying and exploiting the information implicit in large-scale patterns of connectivity[3].

So far, then, I hope to have shown two things. First, that there is enough similarity between the abundance and relevance problems in web-based search and certain aspects of the frame problem to make it worthwhile considering a solution to the former (KP) as at least suggestive of (part of) a solution to the latter. And second, that although KP is designed with a symbolic processing environment in mind, it is possible to imagine connectionist-style versions of the strategy: versions which would combine the idea of gating and control neurons with the use of higher order link-structure information. In the next section, I turn to the 2nd and 3rd worries outlined at the start of this section, viz the balance between search and evaluation, and the threat of circularity.

4. Search, Circularity, and Analogical Reasoning.

The problem of global abductive inference (and, more generally, the frame problem itself) comprises two distinct elements. There is an element of search: the system must find and access the bodies of stored knowledge which are most germane to the task in hand. And there is an element of evaluation: the system must decide how to use, amend, weight or value the information thus retrieved. One possible worry about the use of second-order link-structure information to access and identify the most relevant bodies of stored knowledge is that it leaves this latter problem unresolved.

A simple response would be to accept that something like KP can help, at most, with the search element of the (global abductive reasoning version of) the frame problem. That would itself be a non-trivial contribution. But in fact, I think things look a little better. A lot depends, however, on exactly how we conceive of the remaining problem. Fodor (1983, p. 101-119) distinguishes the two aspects of the problem by analogy with two features of the nondemonstrative fixation of belief in science. The first feature, corresponding to the issue of search, is dubbed “isotropy”. Belief fixation is isotropic if the relevant facts may come from anywhere in a global knowledge-base. In science, Fodor suggests, anything may turn out to be relevant to anything (bacteria to sunspots, etc.). And similarly in individual belief fixation: you cannot know in advance what is relevant to what, hence the need to search, hence the frame problem. The second feature, corresponding to the issue of evaluation, is being "Quinean". In a Quinean system, the degree of confirmation of a belief is sensitive to properties of the entire belief system. Fodor's favorite example (see, e.g. (1983) p. 108) is simplicity. We prefer belief X over belief Y because X yields a simpler account. But simplicity here cannot be measured by any intrinsic properties of X, or by any properties of its encoding or expression (how many sentences it takes to state X, etc.). Instead, we call X simpler if it causes least disruption to, and makes most sense of, the overall set of things we believe. But to know *that* we must somehow be looking at, or sensitive to, the whole belief system. Another daunting, frame-problem redolent task.

Fodor seems to imagine that the second problem (the evaluation problem) requires us to *again* confront the entire knowledge-base. I am not *positive* that Fodor thinks this, but it certainly looks that way. He writes, for example that "the shape of our whole science bears on the epistemic status of each scientific hypothesis" (op cit, p. 107). So one might expect the shape of my whole belief system to bear on my evaluation of the information isolated by KP, or by any other mechanism that successfully confronts the problem of global search.

But this seems wrong. Fodor himself accepts that there is "considerable evidence"[4] that "potentially relevant considerations are often systematically ignored…in favor of relatively local…strategies" (op cit, p. 116). One way to explain such evidence is this: we deploy a moderately reliable procedure to access a delimited subset of knowledge in a problem-solving situation (I have suggested that KP simply reveals a crucial “missing trick” hereabouts). But subsequent evaluation, amendment and action-choice, etc. is then sensitive only to the knowledge thus accessed. The evaluation problem, thus construed, is not to be solved by again confronting the entire knowledge base. And in the *delimited* case, it is possible (as Fodor allows) to imagine various heuristic devices for assessing simplicity, elegance, etc. Fodor’s mistake (see his (1983) p. 11) is to depict the evaluation options as exhausted by 1) the confrontation (again) of the entire knowledge base or 2) the confrontation of an “*arbitrarily delimited*” (op cit, p. 111, emphasis in original) subset of knowledge. The reason a procedure like KP can help with both the search *and* evaluation problems is because it delivers a *non-arbitrary subset* of data over which to define simpler syntactic measures of simplicity, elegance, etc..

In discussing these issues, Fodor considers and rejects one way of delimiting the set of data viz the appeal to something like “frames” or “scripts”[5]. The hope, which Fodor finds forlorn, was that such devices might help by “placing a frame around the body of information that gets called when a given sort of problem is encountered” (op cit, p. 116). But the trouble, as we noted earlier, is that either such frames and scripts are all pre-fixed, ready for every possible problem, or the relevant frames needed to be assembled on the hoof. The former looks impossible while the latter simply recapitulates the problem of identifying the relevant information at the right time. What results is, as Fodor says, a relocation of the problem as one of “executive control” (op cit, p. 117). But we can now see that the contribution of KP is, precisely, a solution to that relocated problem. KP provides a mechanism which allows you to build, on-the-hoof, a temporary module over which to define local processes of reasoning and evaluation. It creates the “*unstable, instantaneous* connectivity” (op cit, p. 118) which allows local processes to access, evaluate and amend just the knowledge structures most relevant to the task at hand. And it does so without itself re-encountering the frame problem. Fodor's general worry, about heuristic-based means of delimiting the set of resources to consider, is that it may require global abductive inference of the same problematic kind in order to decide *which* local heuristics to employ (see Fodor (Ms) p. 41). KP neatly sidesteps this problem, by providing a universally appropriate fast and frugal heuristic viz the identification of hubs and authorities via the use of link-structure information. What KP offers is, I believe, a kind of general-purpose, inward-looking "fast and frugal heuristic" to join the potent group of specialized heuristics treated in e.g. Gigerenzer and Todd (1999).

I turn now to the third and final problem mentioned earlier: the threat of circularity in any attempt to invoke a KP-like procedure in the explanation of individual human reasoning. Kleinberg himself notes that the reason attention to the hyperlink structure patterns among web pages can help solve the abundance problem is that "hyperlinks encode a considerable amount of latent human judgment" and thus carry implicit information about authoritative sources (Kleinberg (1999) p. 2). This may sound worrying. If it takes frame-problem solving humans to set up the hyperlink patterns that KP relies on, don't we need a frame-problem solving homunculus in the head to *set up* the connectivity patterns which would then likewise carry (in the brain) implicit information about what is most relevant to what? The answer is no, we don't. Here's why.

Consider once again the case of hyperlinks in the WWW. And notice, first, that no single designer put that hyperlink structure in place. Instead, there are many thousands of individuals, each pursuing somewhat different local projects, who find it (each for their own local reasons) useful to set up links. Some of those individuals find it useful to set up links to sites with *better* links to sites of common interest. And certain sites will thus be both heavily linked *to* other sites and *from* other sites. This leads to the emergence of what Kleinberg calls "hubs". Moreover, the reason a site becomes a hub is typically because it is linked to many truly authoritative sites. Hence the co-definition, in Kleinberg's system, of hubs and authorities. What I want to stress is that despite the role of human thought in setting up individual hyperlinks, the overall patterns that KP exploits have emerged, even in the web case, as a result of a (globally) blind, self-organizing process. In each case the links are created by thoroughly local processes, none of which itself has access to the overall pattern of events. What KP then does is to exploit information left implicit in such emergent structuring so as to better access the subset of knowledge-structures most likely to be relevant to a specific situation.

The worry about circularity is thus ungrounded. The link-structures emerge, in both web and brain, as a result of local decisions made on the basis of purely local interests and knowledge. What results, however, is a larger pattern of connectivity which is itself the repository of valuable search-sculpting information. The key to effectively searching the web, and (if I am right) to solving a significant fraction of the frame problem itself, lies in the efficient exploitation of this implicit information.

This same link-structure information can, Kleinberg notes, be used in other ways. A natural extension of KP, implemented and tested by Kleinberg, uses link-structure information to discover web pages that are *similar to* some target page. In this case, there is no explicit query-string which can be used to identify the root set of pages from which to trace the in and out degrees of links. It is possible, however, to use KP to identify the strongest authorities in the local link structures leading into and out of the target page. One amends the root-set generating operation and proceeds as before – specifically, the operation of identifying a set number ò of pages containing a target syntactic item (the original query string) is replaced by an operation that finds ò pages that point to the target page. With this foot in the hyperlink-structure door, KP can go on to find hubs and authorities in the *local region* of the target page. Here are the results obtained when the target page (the page such that you are seeking "similar" pages) is www.honda.com:

(www.honda.com) Authorities

.202 http://www.toyota.com/	Welcome to @Toyota
.199 http://www.honda.com/	Honda
.192 http://www.ford.com	Ford Motor Company
.173 http://www.bmwusa.com/	BMW of North America, Inc.
.162 http:// www.volvocars.com/	VOLVO
.158 http:// www.saturncars.com/	Welcome to the Saturn Web Site
.155 http://www.nissanmotors.com/	NISSAN – ENJOY THE RIDE
.145 http://www.audi.com/	Audi Homepage
.139 http://www.4adodge.com/	1997 Dodge Site
.136 http://www.chryslercars.com/	Welcome to Chrysler

To fully appreciate what has here been achieved, reflect that most of these firms will deliberately not mention each other by name or include direct links to each others pages. Indeed, as Kleinberg (op cit, p. 14) notes, many of these pages contain very little text (mostly images) and the text they do contain shows very little overlap. Yet the large hyperlink structure (of links into and out of the pages) contains enough pattern to allow KP, without ever accessing any of the *content* of the pages, to spot the similarities.

It is tempting to speculate (and I hereby give in to temptation) that the close relation between the usefulness of KP in solving a frame-problem related global search problem and its usefulness in driving a process of similarity-based retrieval is far from accidental. For the identification of deep similarities between bodies of knowledge that currently lack short, direct interlinkages looks central to the more general process of analogical reasoning. And analogical reasoning, as Fodor likes to remind us, often seems to exhibit:

> isotropy in the purest form: a process which depends precisely upon the transfer of information among cognitive domains previously assumed to be mutually irrelevant.
>
> Fodor (1983) p. 107

Concerning analogical reasoning, Fodor then asserts that “nobody knows anything about how it works” (op cit, p. 107) and thence proceeds quickly to the infamous First Law of the non-existence of Cognitive Science (see Introduction, above). Although I think Fodor is quite wrong to assert that Cognitive Science has *no* good ideas about analogical reason, I do suspect that KP offers a useful new insight, by once again drawing attention to the possible role of second-order (link structure) information in informing a search for “similar items”.

Other possible extensions of KP include the use of additional kinds of second-order information. (I am calling information second-order when it is not information about the contents of pages or knowledge-structures, but about features of the content-bearers themselves – in particular, for KP the way the pages are linked together). Other kinds of second-order information that might be invoked thus include information about “patterns of traffic” (Kleinberg, op cit, p. 28) ie the way users are moving around the hyperlinked resources, rather than the shape of the links themselves. In the neural case, these may be closely related, with patterns of traffic (propagations of activity) actually creating and modifying the link structures themselves. In fact, KP may be just scraping the surface of a large realm of types and uses of such second-order information.

5. Conclusions: A Second-Order Path From Associations to Reason?

At rock bottom, Fodor ((1983) (Ms)) is worried about three things. He is worried that commonsense reason and abductive inference seems to depend on currently intractable searches of our entire belief-set, or "background of epistemic commitments" (Ms, p. 38). He is worried that sensitivity to local syntactic features is thus often quite insufficient for reasoned response (op cit). And he is worried that appeals to search-reducing heuristics and to associative links between knowledge items at best simply re-position the problem (which then becomes one of knowing *which* heuristics to use, or *which* associative links to follow), again posing an apparently intractable problem of search. That these problems are so closely related as to be pretty much of a piece is admitted on all sides. The common core is the intractability of global search, and it is this intractability that is celebrated in both the frame problem and in Fodor's first law of the non-existence of Cognitive Science.

But the seeds of an algorithmically tractable solution now exist, or so I have claimed, in the (perhaps unlikely) form of a powerful new procedure (Kleinberg, 1999) for searching the web. The procedure (here called KP) sidesteps, whenever possible, the attempt to consult the actual contents of web pages (or stored knowledge structures). Instead, it is driven mainly by second-order information concerning patterns of hyperlinks. The general moral of KP – that information merely implicit in the actual link-structures can be put to valuable work – is, I argued, as applicable to models of individual reason as it is to the web. Nor is it restricted to symbolic models of individual reason – in fact, it comports quite well with recent neuro-connectionist work on “gating” and “routing”, by suggesting an additional layer of information that might be used to fluidly regulate and switch information flow within a complex web of neural populations.

The proposal, viewed in the broadest possible way, is thus to bootstrap an associative engine into a full-blooded engine of reason, not by adding new *kinds* of resources but by better exploiting the information already contained in the associative links. KP, thus viewed, shows how to exploit second-order information concerning patterns of linkage so as to cheaply identify the knowledge-structures most likely to be relevant given a specific input. It is precisely this kind of trick that the developmentalist Annette Karmiloff-Smith (see footnote 5) has in mind when she depicts human thought as the beneficiary of an especially effective procedure for making increasingly efficient use of an original array of merely implicit information. But whereas Karmiloff-Smith and others (myself included) have previously focused attention on the contents of the knowledge-structures themselves, KP invites us to look at the information implicit in the *organization* of those structures. This organization is, I have stressed, the result of multiple locally-based ‘decisions’, and is created without the intervention of any central executive. The proposal is thus that information left implicit by an ongoing process of distributed self-organization may need to be targeted and exploited by other (local, self-organizing) processes so as to turbo-charge processes of reason, recall and adaptive response.

KP is also an example of technology recapitulating phylogeny: a case where the development of web-based search first raises, then begins to solve, a problem deeply analogous to one first encountered by biological brains in the course of cognitive evolution. Fodor's pessimism hereabouts is famous: concerning global search and reason "cognitive science hasn't even *started*; we are literally no further advanced than we were in the darkest days of behaviorism" (1983, p. 129). But look yonder – is that just a web browser shimmering softly, or a glimmer of light in the hyperlinked gloom?

---

[1] One page in R is allowed to bring only some set number of new pages into view, and links between pages with the same domain name ("the first level in the URL string associated with a page") are ignored, as these typically serve merely "navigational" functions within one larger document. The only links that are of interest are thus what Kleinberg calls "transverse links" – links between pages with different domain names. These and a few additional heuristics are described in Kleinberg (1999) p. 6-7.

[2] Suppose that this is wrong, and that the brain (in the kinds of case Fodor considers) uses sentence-like symbolic encodings? In that case, the applicability of something like KP seems evenless problematic. In making the connectionist assumption I am thus making my task much harder.

[3] What is needed is, in fact, something closely akin to what Anette Karmiloff-Smith (1986) (1992) terms a mechanism of "representational redescription". But whereas Karmiloff-Smith imagines that what gets re-described is an actual knowledge-structure, and that these get redescribed in ways that make explicit what was previously implicit *in the knowledge-structure*, what I am imagining is a mechanism which makes explicit the information that was previously implicit in whole patterns of connectivity *between* knowledge-structures. (In the connection setting, in which knowledge is in any case encoded in patterns of weighed connections, this is actually a difference more of focus than of kind). Now imagine a neural network which targets not the outside world, but the activities of other neural networks. What the net cares about, let's assume, is the pattern of temporal relations between activity patterns in different neuronal populations. Over time, such a network will come to encode a kind of low-dimensional image of large-scale patterns of connectivity between knowledge-structures in the brain – it will learn to predict, for example, that activation in such-and-such a population typically leads to activation in such-and-such other areas. Such knowledge might allow such networks to learn to perform the kinds of gating function Van Essen et al describe, in a way that now speaks directly to the search-based aspect of the frame problem itself. Notice, finally, that it would be targeted on the higher-order feature of inner connectivity, and its output would open and close channels so as to create the kind of instantaneous module described in the text. Thornton (2000) describes an approach to "supercharged" learning that looks potentially relevant to this general project.

[4] Fodor cites Nisbett and Ross (1980), and Kuhn (1970)

[5] A script, for example, encodes a body of information concerning the stereotypical course of events in some daily situation (e.g. a visit to a restaurant) – see Schank and Abelson (1975). For frames, see Minsky (1975).