On the Fence: The Role of Language in Cognition

Javier Gomez-Lavin

What is the relationship between language and cognition? Some philosophers argue that language use requires a unique form of mental representation different from those forms used by non-linguistic animals. Others argue that it does not fundamentally alter or restructure the cognitive systems that employ it.

This later view is popular among those that endorse an embedded and embodied approach towards cognition. Embedded and embodied cognition (hereafter referred to as "EEC") replaces complicated internal representations, or mentally instantiated concepts[1], with a continuous interaction between the brain, body and the environment. Take long-division as an example: we rarely do long-division in our heads, relying instead on the interaction between arithmetic thought, pencil and paper. In this example the marks on the page serve as a "scaffold" that simply extends the biologically basic onboard representational systems, allowing us to solve the task with relative ease.[2] It is easy to see how proponents of this approach view language as just such a scaffold – as a provisional, malleable, external structure that extends cognitive processes *without* transforming them.

The most prominent defender of this embedded and embodied approach to language is Andy Clark. In support of his argument, Clark offers a lengthy series of empirical studies. I argue that these studies, when closely examined, do not support his view, and indeed actually lend credibility to the opposing thesis: that language does require a

Chrestomathy: Annual Review of Undergraduate Research, School of Humanities and Social Sciences, School of Languages, Cultures, and World Affairs, College of Charleston
Volume 9 (2010): 64-85

fundamental transformation of the cognitive system. In addition I present one other empirical study that further challenges Clark's account of language and cognition. The goal of this analysis is to motivate an intermediate account of language and its constituent role in cognition, one that neither rejects the overarching aims of EEC, nor eliminates internal representations wholesale.

In the second section I provide a more detailed characterization of the two competing theses of cognition, Clark's version of the EEC and the traditionalist approach. I then focus on each account's take on language, finally sketching out the beginnings of an intermediate compromise. The third section is composed of an extended examination of a few salient studies Clark uses in support of his thesis. I summarize each study, describe Clark's take on the results and, subsequently, demonstrate how they fail to support his avowed view. In addition, I present the results of a supplemental study which motivates the intermediate account introduced in the next section. The final section recounts the evidence provided for this compromise and places it in the context of the larger debate surrounding this divisive issue.

Background

For hundreds of years traditional approaches to understanding cognition conceived of cognition apart from sensation and action. This view originates with Rene Descartes' famous argument that the mind can exist as a distinct entity separate from the body. In Descartes' view the "mind" is conceived as pure thought, understood as distinct from the cognitive processes that are related to the body, such as sensation (Chalmers 15-17). For Descartes the body transmits information, via the senses, to the mind which then processes the information sending instructions directing the body to act accordingly. In the philosophy of artificial intelligence ("AI") this format is known as the "sense-plan-act" architecture (hereafter, "SPA"). The SPA is one way to functionally decompose human behavior (Bechtel 53-62). Relating it back to Descartes' model, sensory inputs transmit information about the environment, encoding it into internal representations. A central processor, the "mind," then assesses the information and forms a plan, which is relayed to the body. Finally, the body carries out the plan by, for example, moving about or manipulating the environment as directed.

Think of this process in the game of chess. In a typical chess game, a player will first look at the board, noticing all of the pieces along with their distinct positions. In doing this the player gathers information about the current state of the game, e.g., whether a king is in check or a pawn threatened – this is equivalent to the "sense" part of SPA. Given that the player knows the rules of the game, she will then combine this information with the proscribed rules to form a "map" of her possible moves. In weighing the consequences of her possible moves she will attempt to select the move with the most beneficial outcome, e.g., sacrificing a pawn to remove the king from danger. This cognitive process of managing and manipulating information is the "planning" in SPA. Lastly, the player physically manipulates the piece to carry out the planned move; this is, not surprisingly, the final part of the triptych – action. It is important to note that the entirety of the cognitive (thought) process takes part *internally* and *separately* from the player's perception of and action on the environment.

In order to enact the SPA, the agent in question needs to possess a full store of internal "on-board" mental representations. For example, our chess player needs to internally represent the state of the chess-game at any given time. We can imagine this as a "map" with distinct representations containing information particular to, and standing in for, any given piece. In order to effectively play the game, the player needs at least two different internal representations. The first is a representation of the current state of the board, or a mental map of the arrangement of the pieces. The second is a representation of the rules of the game: the allowed moves of each piece, the endpoint or goal of the game, etc. From the SPA account, it is almost as if our player is internally recapitulating the properties and states of the physical entities as they exist out in the world.

Many philosophers have criticized the SPA paradigm in AI and cognitive science from its inception (Bechtel 62).[3] However, it was only in the latter part of the twentieth-century, from the 1980's onwards, that these criticisms were developed into tenable theses opposing the traditionalist view. The EEC account of cognition is just such a thesis. As prefaced above, proponents of EEC suggest that one *cannot* separate cognition from sensation and action. In essence, cognition doesn't happen in a vacuum; one cannot successfully disconnect cognition – the active managing and processing of information – from its context, its environment.

Unlike SPA, the EEC account takes cognition to be an active and continuous dialogue between the "seat" of information-processing (i.e., the brain) and the world around it (Clark 5-10). One could envision it as a feedback loop or a reflexive process: the cognitive agent manipulates the environment and uses the outcome to reassess its situation and so on. Consider someone walking into a dimly lit warehouse they have never been in before. Following the SPA paradigm, our adventurous cognitive agent would: (i) survey the expansive warehouse, (ii) form a mental "map" with representations of the various items scattered about (e.g., huge stacks of crates, forklifts, pallets, barrels, machinery etc.), (iii) envision a plan to navigate the expansive maze-like warehouse, and finally (iv) carry out the plan. However, this seems impractical if not unnatural; one doesn't stand at the front of a warehouse attempting to compute all the variables, routes and possible actions available for every situation. This is especially evidenced in cases where actions are temporally constrained, for example, if the warehouse were on fire. A more natural course of action for our adventurer would be to move through the warehouse focusing on discreet and salient bits of information (e.g., more light coming from one side of the building, less smoke coming from another etc.) and using this information to constantly reassess his situation and change his course through the maze. By focusing only on the most salient pieces of information as opposed to reconstructing the whole environment as an onboard internal map, the agent is able to efficiently navigate his world.

Empirically conducted studies have lent added credibility to the EEC approach. For example, David Kirsh and Paul Maglio studied how humans go about playing the game Tetris; for a detailed discussion see Clark and Chalmers (10-12) and Clark (48). In Tetris, the player is confronted with several possible geometric shapes; some are squares, others are lines, and yet others "Z" or "T" shaped etc. These shapes fall towards the bottom of the screen, and the player's job is to rotate them so that they form interlocking rows. Once a row is complete, it disappears; incomplete rows will remain and eventually fill the screen, resulting in the game ending.

If our cognition resembled a traditional SPA architecture one would expect that a player would (i) perceive the shape, (ii) mentally rotate it so that it fits in with the unfilled row below, and then (iii) carry out the rotation by pressing the appropriate button. In this case, all of the rotation occurs in the central processor, or in the head of the player. In contrast, Kirsh and Maglio found that the player physically rotates the shape using the buttons and then compares the result of the transformation with the unfilled row below. When a rotation is found that appears to work, the player drops it into place. Their study demonstrated the effectiveness of this approach: it turns out that it takes approximately 1000 milliseconds (ms) to mentally rotate a shape while only 300 ms to perform a physical 90° rotation (Clark 71-73). The crucial point is that the planning does not take place in the central processor in the head, but rather in a feedback loop combining the environment, sensory inputs (e.g., vision), onboard information processing systems (e.g., visual-spatial processing in particular regions of the brain) and actuators (e.g., innervated muscles, hands and fingers).

The ability to physically, as opposed to mentally, rotate the shape allows us to take some of the burden off of our cognitive system. Similar to the long-division example above, our manipulation of the world around us – whether by pressing a button to rotate shapes, or writing out the steps of a division problem – removes some of the difficulty, thus allowing us to extend our cognitive capacity (Clark 72-73). The action of rotating the shape does not, in and of itself, get the player closer to the "goal," i.e., filling the rows. Rather, it gives the player more information about the relationship between the outline of the shape and the unfilled row at the bottom. This information then allows us to fill the rows, or solve the problem, or reach a particular goal *faster* and with greater *efficacy*. The manipulation of our environment prepares us for efficient and pragmatic performance.

This notion may seem a bit foreign at first, but we manipulate our world to help us solve tasks every day. Take cooking as a quick example: people often spatially organize the recipe they are going to prepare, wet ingredients on the left, dry ingredients on the right; produce in the bowl, uncooked meat in the Tupperware, etc. We place the dog's leash on the door handle as a reminder to take him out on a walk. We stick a note on the refrigerator reminding us to get milk and eggs. In effect, all we are achieving is a reorganization of the world around us, which helps mitigate our cognitive burden and makes our lives easier.

By describing this process more accurately, EEC has radically changed how we view and study cognition. It frees us from a robotic paradigm of perceive → compute → execute. It forces us to take the environment into account when studying cognition or AI. If an EEC account of cognition is correct, then we should not think of cognition as necessarily recapitulating, or transcribing, features of the world into the head. We do not have to conceive of cognition as importing a map or model of the Tetris board, chess board or warehouse; instead, we can conceive of cognitive agents simply referring to it, when needed, as it exists out in the world. Following EEC, cognitive systems let the world serve as its *own* best model, tweaking and cajoling it when needed to help us solve the task(s) at hand. The debate now turns on perhaps the most human trait of all: language.

Language is a highly structured feature of our environment. Our environment is saturated with linguistic entities in the form of written and spoken sentences, and those sentences have two primary constituents: individual words, and the rules – the grammar – we use to combine them. The constituent structure of language is its grammar, or the syntax used to combine individual meaningful pieces (i.e., words) into longer strings (i.e., sentences) that compose a language. For example, in our language, English, the structure prescribes how to fit words of different classes together: nouns with adjectives, verbs with adverbs, subjects with predicates, etc. This structure allows us to form and communicate an infinite number of correct (i.e., grammatical) strings (i.e., sentences) from a finite storehouse of individual, atomic, components (i.e., words). Without this combinatorial structure guiding language we would have to have an infinite number of novel templates each corresponding to particular expressed thought. An example of this combinatorial structure is the phrase: "John loves Mary." "John loves Mary" follows the prototypical subject-verb-object structure, while using three finite elements: John, loves and Mary. With just a few elements, or words, in play, the combinatorial structure allows me to form several different grammatical (and meaningful) strings: John loves Mary; Mary loves John; John and Mary love X; X loves John and Mary; X loves John not Mary; X loves Mary not John, etc. If language did not follow a consistent structure we would have to add additional unstructured (grammatically incoherent) strings to represent other combinations of words.

As one might foresee from the discussion above, defenders of EEC would like to leave language and its structure out in the world. In particular, for Clark, language should be viewed in the same way the EEC treats chess, long division, and Tetris. In the same way that manipulating a shape in Tetris can reveal a solution to the problem of fitting the shape into an unfinished row, linguistic structures can redirect our onboard representational resources in novel ways (45). Recall that we often manipulate our physical environment to make our lives easier: organizing the pantry so that the foods we consume the most are easily accessible, placing your car keys next to your wallet etc. According to Clark, we use words in the same way. They extend our ability to organize ideas an extraordinary degree. Words, along with their constituent structure, enable us to organize hugely disparate concepts such as "Juliet" and "the Sun" in the same meaningful space, as Romeo does when he says, "But soft! What light through yonder window breaks? / It is the East, and Juliet is the sun!" (II.1.44-45). Language, for the EEC proponent, gives rise to vastly more nuanced and complex forms of environmental manipulation or reorganization than were possible before it existed.

Perhaps that gloss is a bit too simplistic. Clark notes that language provides three facets that come together to form a "scaffold" for the cognitive system to act on. For Clark, language affords access to a cheap, arbitrary, context-free, labeling mechanism that allows the cognitive system to divvy its environment into finer or larger-grained categories. It also encodes behaviors which – when presented to or recalled by the cognitive system – allow the cognitive system to gain a familiarity and expertise not attainable by other means.[4] Finally, Clark believes that language provides a unique toolset enabling the cognitive system to engage in meta-cognition or thinking about thinking, thus giving cognitive systems insight and some control over their own patterns of thought (44).

Clark's account of language provides a unique, new explanation of "material symbols" (e.g., written or spoken words and numerals) and their coded strings (e.g., sentences), showing us how they enable disparate modality-specific processing regions of the brain (e.g., visual, auditory, tactical etc.) to perform complex and novel tasks. Furthermore, and most importantly, Clark argues that the implicit structure of language, as described above, needn't be replicated in the head. For example, he states:

> Words and sentences, on this view, may be potent real-world structures (material symbols), many of whose features and properties (arbitrary amodal nature, extreme compactness and abstraction, compositional structure, etc.) *simply complement, without full replication,* the contributions of basic biological cognition. (55, emphasis added)

The question then arises: Is it possible for a cognitive system to use language without a profound transformation or reorganization of internal representational structures, so as to accommodate the structure of language?

The intuition against this perspective is immediately apparent. Consider a key. The key has a certain structure of grooves and ridges. But that structure is irrelevant if there is not a lock whose own organization is sensitive to the structure of that key. Intuitively, sentences are like keys and it seems sensible to think that a corresponding lock may lie somewhere in the cognitive system.[5] From this perspective, the lock possesses corresponding structural information about the key. Without this information, embodied in the tumblers, the key could not fit into, or casually interact by turning, the lock. Much in the same way, it seems that a cognitive agent would be unable to use or manipulate words and sentences, either spoken or printed, without some kind of corresponding information about their structure stored *within* its own cognitive mechanism.

Yet throughout his account of language, Clark refuses to entertain the notion that a cognitive agent has to import the structure of language as prerequisite for its use. When responding to Michael Wheeler's critique of the EEC account of language, Clark writes, "[I]t must be possible to represent syntactically structured language without using syntactically structured representations to do so (just as it is possible to represent green objects without using green representations to do so)" (57). Surely, I can represent the common phrase "the dog bit the boy" without having knowledge of its implicit structure; however, it would probably appear to me as some variation of a mental image (e.g., close your eyes and imagine a house – that would qualify as a mental image of a house). It is as if words, as they exist out in the world, lack the information necessary to be used by a cognitive system that does not have the prerequisite structure with which to process the words. Without this structure, or lock, it becomes apparent that I cannot use the string of words to guide and shape cognitive processes, much less behavior.

Overall, it is clear that Clark is exceedingly cautious in allowing restructuring *in the head* However, as Whitney Schonbein notes, language is a particular case that poses a "unique type of representational problem" (8). Clark must confront the notion that any cognitive system with finite capacity must find an ecologically efficient method to represent the *infinite* number of structured strings of words (e.g., sentences) that constitute natural language. Obviously, as Schonbein points out, this infinite "set cannot be represented by an infinite number of distinct, unstructured surrogates, because our onboard representational resources are limited" (8). So, for the linguistic external scaffold to have its desired unifying and edifying effects on the cognitive system, we may just need to elicit the support of a productive, systematic, parallel, inner scaffold.[6]

It is starting to appear that Clark may have overreached and crafted too strict of a proposal with respect to language's role in cognition. I would like to sketch an intermediate account roughly following the lock and key analogy, and similar to those proposed by Schonbein and Wheeler. This account will grant many of the insights of EEC – such as external scaffolding – while endorsing the claim that linguistic mental representations will respect and recapitulate the structure of language. Before I outline this intermediate account I will consider the empirical evidence cited by Clark in support of his thesis.

Analysis

In my analysis I will dissect two of the studies in which Clark argues that language forms an external scaffold for the cognitive system and, consequently, that it is inessential –if not redundant – to recapitulate this structure in the head. I will then present the results of a newer study performed by Dehaene et al. which contradicts these findings and may better support the intermediary view introduced above. Before I examine the details of the studies, perhaps a brief overview of Clark's general argument is in order:

1. If the EEC approach to language's role in cognitive systems is correct, then linguistic structure (i.e., syntax), – as it exists in the external environment – need not be replicated in the internal representational systems.

2. Given the empirical evidence provided by conceptual matching and arithmetic studies of language supporting EEC, it follows that: (i) the transformational role of language on cognition is due, by in large, to language and linguistic structure acting as an external attention-focusing scaffold, enabling the cognitive system to accomplish novel and complex tasks (e.g., long-division), and (ii) there does not appear to be a need to reorganize the internal representational systems (i.e., to train or to retool the brain), in order for cognitive systems to use linguistically structured representations to accomplish these tasks.

3. Therefore, linguistic structure can be left in the world. Essentially, there is no *need* to translate or reproduce the linguistic structure in the head.

Roughly outlined, we see that having the internal representational systems (often referred to as "vehicles" in the literature) adopt linguistic structure (a combinatorial and productive syntax) is anathema to Clark.

Clark argues that recapitulating the structure of language in the head would be an evolutionarily expensive and demanding pursuit. After all, would it not be redundant, or at least inefficient, to have evolved to recapitulate structure that already exists as complete out in the world? For example, our brains did not evolve to map and compute the exact tension and flexion of each individual muscle fiber in our legs, thus enabling us to walk. Instead our brains rely on the structure inherent in the anatomy of our legs; merely sending a "start" signal to the muscle groups that then contract in reflexive ways (this reflex can often be seen in newborns who consistently and instinctively engage in kicking motions when placed on their back, or in paralyzed individuals who receive mild electrical stimulation in their spine). These innervated muscles then send back sensory information to the brain, precipitating a feedback-loop like process, whereby the brain modifies the signal to increase or decrease walking speed etc. (Clark 3-9). Clark's view simply has linguistic structure mirror this example; allowing the cognitive system to use language without onboard restructuring.

Two of the studies on which Clark relies are Thompson, Oden and Boysen's 1997 paper on conceptual cue matching ability in language-naïve chimpanzees and Dehaene et al.'s 1999 study examining the sources of mathematical thinking through neuroimaging. My goal in revisiting these studies is to demonstrate how their results actually weaken both the second half of the conditional in premise (1) and the entirety of premise (2). In effect, I will attempt to perform a *modus tollens* on the argument sketched out above. (For those unfamiliar with the procedure, a *modus tollens* is a logical proof that enables the user to deny the consequent, or conclusion, of a particular set of arguments in standard form.) Clark's argument can be stylized as such:

1. If *P* (i.e., if the EEC account of language is correct), then *Q* (i.e., then linguistic structure is not recapitulated).

2. *P* (i.e., empirical studies support the EEC account of language in certain ways).

3. 3) Therefore *Q* (i.e., linguistic structure is not recapitulated).

A *modus tollens* will first show that *Q* is not the case, thereby disproving *P*. So, a sketch of my argument would be as follows:

1. If *P* (i.e., if the EEC account of language is correct), then *Q* (i.e., then linguistic structure is not recapitulated) – as discussed above.

2. But, in fact, it is not the case that *Q*. The empirical studies discussed below fail to show that linguistic structure is not recapitulated as a prerequisite for language-use.

3. Therefore it is not the case that *P* (i.e., the EEC account of language is *in*correct).

Three studies

In "Language-Naive Chimpanzees (*Pan troglodytes*) Judge Relations Between Relations in a Conceptual Matching-to-Sample Task," Thompson, Oden and Boysen attempt to disprove previous conceptual matching task (CMT) research, which suggests that only language-trained (LT) chimpanzees are capable of spontaneously matching abstract relations between novel pairs of objects. The CMT, when simplified, goes something like this: a set of artifacts exhibiting a first-order "sameness-relation" (such as two green cans, or two blue shoes) are presented to the subject. We can abstract these objects as sets of letters, e.g., AA or BB; however, it is important to note that the subjects are not shown a set of letters, rather only the aforementioned set of objects. Other objects presented to the subject could exhibit a first-order "difference-relation" (such as a yellow bottle and a green can) that we can abstract as CD or EF. The difficultly comes in play when the subjects are prompted to match groups of the pairings for sameness or difference: e.g., AA and BB as well as CD and EF would have a second-order "sameness" relation as both pairs are composed of objects of a similar first-order relation. Furthermore, AA and CD or BB and EF would be matched as different, as the instantiated first-order relation between the pairs is different; one has a first-order "sameness" while the second has a first-order difference.

The CMT literature suggests that only chimpanzees extensively trained in language and numeral acquisition, or able to put together compositional strings of symbolic tokens according to predetermined syntactical rules (a grammar), would be capable of picking up the second-order task with ease (Thompson et al. 34). Previous research understood that the chimpanzees, with their representational system now reformatted by the acquisition of linguistic structure, would quickly conceptualize the sameness or difference relations between the pairs. Indeed, the previous literature only showed two ways around the problem: You either trained the chimpanzees in language-like problem solving behavior or you performed "dogged" reinforcement; essentially, taking a language-naïve (LN) chimpanzee and subjecting it to thousands of trials and routines of reinforcement, thus training it to successfully match concrete relationships of a certain type (e.g., ½ an apple is "the same as" ½ a glass of water) (41).

Thompson, Oden and Boysen aquired one LT and four LN chimpanzees and trained them in token recognition. This is very different than comprehensive language training, as there is no syntactical or combinatorial element learned. Rather, when presented with a sameness relation, AA, they were also presented with a heart-shaped plastic token. Likewise, when presented with a difference relation, EF, they were presented with a diamond-shaped token. After token training, the subjects were then prompted to perform the second-order CMT described above. It is important to note that the tokens were not physically represented along with the pairs of objects. Still, their results demonstrated that after only a few rounds (about 32 to 64 matching exposures) four of the five subjects successfully and consistently matched these second-order relations far above what could be predicted by chance (on average above 80% correct) (37).[7] The few rounds of exposure are substantially less than the predicted thousands of repeated exposures thought to be necessary for LN chimpanzees to master the CMT. Also of interest was the fact that three of the four LN chimpanzees performed statistically identically to the one LT chimpanzee.

The experimenters then generalized their findings to suggest that the token training, which also happens to be the foundation for more advance language training in chimpanzees, forms the lowest-common denominator between the two successful groups (42). Therefore, it is not competence with syntactical structures that predicts successful second-order conceptual matching, but rather familiarity with the arbitrary tokens which stand in place of the original first-order relations. Essentially, when presented with an AA and BB pairing, the subjects virtually tokened two heart shaped symbols, thus collapsing the size of the task at hand to the vastly simpler first-order relation. Hence, no real restructuring or reformatting of the biologically basic onboard representational systems to accommodate linguistic-like structure (i.e., syntax) is necessary for the subjects to perform the CMT.

At first glance this thought provoking study seems to support Clark's claim that material symbols (plastic tokens in this case) serve as external prompts which focus the onboard, evolutionarily traceable, biologically basic systems of attention to salient features and patterns in the organism's environment. Furthermore, it seems apparent that the subjects in this study do not recapitulate or translate the tokens into atomic pieces which serve their functional role in a syntactically structured system; the tokens merely serve as surrogates that group features of the environment onto a new virtual space. However, it seems like a stretch to follow Clark's conclusion that language must then leave its structure out in the world. I do not suggest that he is pulling the wool over our eyes; however, these tokens *prima facie* do not mirror linguistic structure – *they are unstructured*! The tokens follow no grammar, no syntax, nor a compositional form with which to arrange them to express novel combinations.

I agree with the authors that this type of surrogate tokening is the lowest-common denominator necessary to perform higher order CMTs. Surrogate tokening, in one form or another, does seem to function as a prerequisite to structured language use in general. However to go as far as stating that the "present results… support the theoretical assumption that symbols are 'in the world' first and only later 'in the head'" is a far leap (42). In the CMT, there are no requirements on the cognitive system to use computational, or rule-governed, manipulations of these tokens; they merely stand for a particular relationship in the world. Essentially, unlike language, the tokens lack structure. Consequently, the cognitive system does not require a structured representational system (i.e., a "lock") in order to make use of the tokens.

In the case presented, the tokens do themselves act as an attention-focusing scaffold, by reducing the complexity of the second-order relationship. Yet, language has structure which needs to be paralleled by the onboard representational systems themselves; for without this efficient structuring the cognitive system would be flooded by an infinite quantity of unstructured tokens acting as surrogates for countless relations.[8]

Clark uses the 1999 study by Dehaene et al. to demonstrate that research in arithmetical thought supports the EEC model of the cognitive system as efficient, non-SPA, pattern completing system (Clark 51). In their paper, Dehaene et al. use sophisticated neuroimaging techniques to try to isolate the particular neurological structures implicated in arithmetic thought. Given their results, it will become clear that this study in fact helps support an intermediary position between external and internal structuring.

Here is a rough sketch of their experimental design: Dehaene et al. first took two groups of bilingual Russian-English individuals. They trained both groups in simple arithmetic addition; however, they trained each group in a separate language. The researchers then quizzed participants in exact (e.g., "does $2+4 = 6$ or 8") or approximate (e.g., "does $2+4 \approx 5$ or 9") addition. Their results showed that in the exact addition category the participants responded statistically faster (on the order of 1000-1200ms) when queried in the language in which they were taught ("Sources" 971). The researchers go on to posit that the faster reaction time to the trained problems indicates that "each new fact was stored independently of neighboring magnitudes, perhaps as a sequence of words" ("Sources" 971). However, there was no significant difference in response time when the participants were asked approximate addition – in either the trained or untrained language. These results seem to indicate that exact arithmetical thought might be processed in more language-dependant regions of the brain, whereas approximate addition was processed in language-independent, visuo-spatial regions ("Sources" 972).

When the researchers placed the participants into MRI machines while performing the arithmetic tasks, they found that different regions of the brain showed activity when calculating exact or approximate addition. When responding to exact equations, regions of the left inferior frontal lobe were active. Previous studies have also found "left inferior frontal activation during verbal association tasks, including generating a verb associated with a noun" ("Sources" 972). On the other hand, participants performing approximate addition showed increased activity in the visuo-spatial processing circuits of the dorsal parietal pathway, near the occipital lobe. This also corresponds with previous research done in patients with severe left-hemispheric lesions who "could not decide whether 2+2 was 3 or 4… but consistently preferred 3 over 9" ("Sources" 973). Overall, we see that their results support the notion that even basic arithmetic invokes multiple representations in several regions of the cognitive system.

The overarching conclusion reached by Dehaene et al., that even elementary arithmetic employs the use of multimodal forms of mental representation, does leave the door open for Clark and the EEC paradigm. It also lends support to the notion that the cognitive system is an adaptable, multimodal, processing machine. Furthermore, and most importantly, it is now clear that certain localized brain regions trade in structured representations. While the idea that an arithmetic symbol, for example "98," is translated into an equivalent internal representational element might not be fully supported by Dehaene et al., we do see that there is a place in the brain where the numerals, at least in exact addition, are engaged in a computational –perhaps combinatorial – process, recapitulated and fitted into specific language dependent representational systems which then run their course to compute a solution.[9]

Clark does not completely concede this point. He states that the work of Dehaene et al. is evidence of a "hybridized" representational vehicle that may include content-relevant internal representations along with "a co-opted proper part, a token (let's think of it as an image, very broadly construed) of a conventional public language encoding ('ninety-eight')" (53). Clark seems to be slowly, cautiously, shying away from a complete externalist picture of language; yet, once you open the door to internal representations, it becomes hard to mitigate the effects they might have on the adaptive, malleable and efficient cognitive system. I believe this last study demonstrates that point succinctly.

In a newer paper, "Log or Linear? Distinct Intuitions of the Number Scale in Western and Amazonian Indigene Cultures" (2008), Dehaene et al. attempt to discover whether the linear number-line scale used in western arithmetic thought is a byproduct of an eventual biological maturation or a taught cultural phenomenon. They referenced several studies which point to young human children, infants and other non-human primates representing numerical quantities logarithmically along a number line ("Log" 1217). For example, given a one-dimensional finite line, with zero on the left and 100 on the right, western children consistently represented the smaller quantities towards the left and the larger quantities towards the right ("Log" 1217). However, their representation was disproportional, mirroring a logarithmic scale: So, for example, they often placed 10 in the middle of the line and grouped the larger numbers into a much smaller space than one would on a linear number line. These studies indicate that these children quickly adapt to the linear configuration, often phasing out the logarithmic scale by age 10 ("Log" 1219). Following this notion, Dehaene et al. decided to test how the Munduruсu Amazonian tribe, a group with a reduced numerical lexicon and little education, would format number lines.

The experiment went roughly as follows: The researchers presented 33 Munduruсu adults and children with a digitally simulated horizontal segment that had a movable cursor controlled via a mouse. The left of this line had a white circle with a single black dot present. The right side had either 10 or 100 dots present, depending on the experimental condition (either 1-10 or 1-100, respectively) ("Log" 1217-18). Each participant was then presented with either a separate display of dots numbering between the two polar quantities, a series of tones similarly quantified, a spoken Munduruсu phrase or a spoken Portuguese word ("Log" 1218). They were then asked to place the cursor approximately where they thought the indicated quantity should lie. They then compared these results with data procured from American participants with some form of formal education ("Log" 1217).

Overall, through several sophisticated measures and analyses, Dehaene et al. found that, "Logarithmic thinking persists into adulthood for the Munduruсu, even for very small numbers in the range from 1 to 10, whether presented as dots, tones, or spoken Munduruсu words" ("Log" 1219). This stands in distinct opposition to the data gathered from the American participants; Dehaene et al. cogently deduce that "In light of the performance of Amazonian adults, it is clear that the mental revolution in Western children's number line does not result from a simple maturation process" ("Log" 1219). They then extend their data to support the idea that the concept of linearity, along with a similarly structured number line, is by in large a cultural innovation. Thus, while numerical intuition, mapping numbers onto some spatial domain, seems universal, culturally specific innovations alter this representational map, with several consequences ("Log" 1217).

Further supporting research demonstrates that there may be specific neural circuits whose responsibility it may be to act as mechanisms of numerical perception "whereby individual neurons in the parietal and prefrontal cortex [some of the most developed cortical regions] exhibit a Guassian tuning curve on a logarithmic axis of number" ("Log" 1219). It then appears that evolution may have selected an onboard biologically basic logarithmic scale for its overall representational efficiency; as it is capable of representing several numerical orders of magnitude in a relatively small package (Nieder and Miller 7457). Once linear number lines are introduced, it seems that a reorganization of some of the surrounding structures involved in numero-spatial representation may be in order. However, to obtain definitive proof that some onboard restructuring is taking place we would have to observe a neuronal plasticity induced by arithmetic training in vivo, *in action*. Yet, in light of these findings it now seems plausible to suggest that some form of enculturation or training can reshape or reformat these biologically basic processes to display more demanding and complex forms of linearity and arithmetic. Especially when paired with the 1999 study by Dehaene et al., we can now see that there might just be some internal tacit reorganization or restructuring of the internal representational systems (i.e., the locks) according to structural (in this case, spatial) commitments carried by the internal representations (i.e., the keys) of the numerals themselves.

Synthesis

Given the three studies examined above, we can see how their results chip away at Clark's stubborn refusal to grant any reformatting of a cognitive system's internal representational systems along the lines of linguistic structure. In the first study we learn how tokened representations may very well underlie much of linguistic acquisition; however, we simultaneously see that these tokens are in-of-themselves unstructured and structurally impotent, a critical difference from the material symbols associated with language. In the second study, we find that structured (arithmetic) thought invokes a combination of multimodal representational systems (language dependent and visuo-spatial); however, we also find a particular, localized, process through which representations of material symbols enter into a linguistically structured format. Finally, in the last study we learn about a process that might epitomize how acculturation and learning might come to induce a (linguistically) structured reorganization in particular onboard representational systems underlying cognition.

Convincingly, when taken together, the results point us back to a basic premise of the EEC approach: The importance of an interplay and exchange between a cognitive system's external environment and the malleable representational systems it harbors; systems which are *constantly* changing and adapting to new representations (such as the relatively recent [±50,000 years] rise of language itself) in order to seek out the most efficient way to navigate and manipulate their environments.

This analysis demonstrates both the power of the external scaffold and the impressive adaptability that onboard representational systems of cognition wield. It seeks to map out an intermediary position between the wholesale linguistic structuring of thought and an all-out externalist take on language. It now seems sensible that if the cognitive system can encounter and entertain internal representations of language and material symbols that it would, perhaps through evolutionary pressures, pick up on the structure of language. On the whole, the paradigm of EEC does not contain a clause singling out any-and-all forms of onboard representational restructuring. Why wouldn't the adaptive cognitive system pick up on the infinitely productive, compositional structure of language and mold itself to extract its inherent productivity and efficacy, thus embodying its structure? Considering the sheer length of time on an evolutionary scale this may very well have happened.

Returning to Clark's view, we now see that his refusal to have onboard recapitulation of the structure of language is attacked on two fronts. Driven both by contemporary research in his preferred model of biologically basic cognition, i.e., artificial neural networks as showcased in Schonbein's paper, and by data collected from contemporary scientific studies examining numeral and language use by primates and humans. Through these studies and due to the extensive work of philosophers and psychologists alike, we see that it might just be *impossible* to deny the transformational impact that language has on the cognitive system. However, we must give Clark his due. He may have accounted for this argument, as he cryptically leaves the potential for the blurring of the external/internal dichotomy:

> From sounds in the air to inscriptions on the printed page, the material structures of language both reflect, and then systematically transform, our thinking and reasoning about the world. As a result, our cognitive relation to our own words and

> language (both as individuals and as a species) defies any simple logic of inner versus outer. (59)

As I see it, this statement may undermine some of the claims contained in his book. It is not that Clark waffled mid-way through his work; rather he leaves the possibility open that internal restructuring may end up playing a role in this incredibly complex cognitive narrative.

Perhaps Clark is worried that dissolving this external/internal divide to include comprehensive internal restructuring will result in a linguist's dream –and Clark's nightmare – in which the biological elegance of cognition will be replaced piece by piece with computer-like formal language and its syntactical accomplice. Instead of putting blinders on to the potential of internal restructuring, we should redouble our efforts towards outlining the extent of just such a process. Successful implementation of computational models, such as advanced AIs, might just hold the key to understanding which features, structures and processes we should look for in the head. Armed with this insight and information, we might be better able to direct our resources and research towards understanding those processes and features that underpin cognition. And in the end, isn't that our goal as researchers of cognition?

Notes

Thanks to Dr. Whit Schonbein and Brian Everett for their helpful comments and suggestions.

[1] This positive account of internal, mental, representation is a very broad gloss of much of the philosophical literature. Perhaps the most accessible manner to envision a mental representation is to close one's eyes and imagine a dog. Whatever happens to come up in your mind, be it a rough image of your favorite household pet or a word-tree of concepts related to "dog" (such as cats, Lassie, etc.) is an internal, mentally instantiated, representation of the concept "dog." As you can readily see, there exists substantial room for variation and complexity.

[2] In this case the term "biologically basic onboard representational systems" refers to an agent's innate information processing "hardware," or for our case, as humans, our overall nervous system consisting of our brain and the sensory and perceptual components (eyes, ears,

nerves etc.) that complement it.

[3] Hubert Dreyfus was one of the first such critics, suggesting that the formalized logic and symbolic representational format of AI fails to capture lived-experience, or the experience of being an actual agent in the physical world, a notion philosophers often refer to as *phenomenological experience* (Bechtel 62).

[4] Think of a child tying her shoelaces. Repeating a learned mantra (e.g., "over and under it's easy to do/ reach through the hole and pull it through") helps organize a behavior into manageable pieces, allowing the agent to eventually gain proficiency with the task-at-hand until they no longer have to continually recall the linguistic support-structure. Essentially, by recalling these linguistic supports, an agent self-sculpts behavior, eventually rendering the "scaffold" unnecessary.

[5] I am adapting and extrapolating this lock and key analogy from Hugh Clapin, who uses the analogy with regard to mental content and functional cognitive architectures.

[6] Schonbein develops and continues this line of thought in "The Linguistic Subversion of Mental Representation."

[7] This success rate was consistently found even after controlling for the effects of variable interactions, confounding variables, differing reinforcement schedules, and potential unintentional experimenter influence (Thompson et al. 34).

[8] What may be an interesting next step is to see if these LN chimpanzees, once familiarized with tokens, would be able to perform computational "third-order" or "forth-order" relations. For example, if presented with two differing second-order relations: ◊&♥ and ◊&♥, then the chimpanzee would have to select, on a touch-screen perhaps, a third symbol: ¤, as opposed to merely identifying the third-order relation with the second-order sameness token. Likewise, when presented with ◊&◊ + ♥&♥ the chimpanzee would select: §. The results garnered from such a study might better support the claim that cognitive systems can actually *utilize* material symbols without formal training or the transformation of internal representational systems

[9] Recall that the subjects trained in one language had significant response related deficits when switching over into the untrained language; hence this region must be very closely tied into language-dependant structures.

Works Cited

Bechtel, William, Adele Abrahamsen and George Graham. "The Life of Cognitive Science." *A Companion to Cognitive Science*. Ed. William Bechtel and George Graham. Oxford: Blackwell, 1998. 1-104. Print.

Chalmers, David J. "Rene Descartes: Meditations on First Philosophy." *Philosophy of Mind: Classical and Contemporary Readings*. Oxford: Oxford UP, 2002. 10-20. Print.

Clapin, Hugh, ed. *Philosophy of Mental Representation*. Oxford: Oxford UP, 2002. 295-97. Print.

Clark, Andy and David Chalmers. "The Extended Mind." *Analysis* 58 (1998): 10-23. Print.

Clark, Andy. *Supersizing the Mind*. Oxford: Oxford UP, 2008. Print.

Dehaene, S., et al. "Log or Linear? Distinct Intuitions of the Number Scale in Western and Amazonian Indigene Cultures." *Science* 320 (2008): 1217-20. Print.

—. "Sources of Mathematical Thinking: Behavioral and Brain-Imaging Evidence." *Science* 284.5416 (1999): 970-74. Print.

Nieder, A. and E. K. Miller. "A Parieto-frontal Network for Visual Numerical Information in the Monkey." *Proceedings of the National Academy of Sciences USA* 101 (2004): 7457. Print.

Schonbein, W. Whitney. "The Linguistic Subversion of Mental Representation." The College of Charleston, 2010. TS.

Shakespeare, William. *Romeo and Juliet*. *The Norton Shakespeare*. Ed. Stephen Greenblatt. Vol. I. New York: Norton, 2008. 897-972. Print.

Thompson, Roger K. R., David L. Oden and Sarah T. Boysen. "Language-Naive Chimpanzees (Pan troglodytes) Judge Relations between Relations in a Conceptual Matching-to-Sample Task." *Journal of Experimental Psychology: Animal Behavior Processes* 23.1 (1997): 31-42. Print.

Wheeler, M. "Is Language the Ultimate Artifact?" *Language Sciences* 26.6 (2004). 693-714. Print.