Title
From Neurons to Culture: Applying Newell's Systems Levels to Understanding the Impact of Culture

Journal
Proceedings of the Annual Meeting of the Cognitive Science Society, 44(44)

Authors
West, Robert L
Shaeen, Tristan A.

Publication Date
2022

Peer reviewed

From Neurons to Culture:
Applying Newell's Systems Levels to Understanding the Impact of Culture

Robert L. West (robert.west@carleton.ca)
Department of Cognitive Science, 1125 Colonel By Drive
Ottawa, ON K1S 5B6 CA

Tristan A. Shaeen (tristan.shaeen@carleton.ca)
Department of Cognitive Science, 1125 Colonel By Drive
Ottawa, ON K1S 5B6 CA

Abstract

Allen Newell laid down highly influential principles for the study of cognition with his systems levels framework, but it is less well known that this framework also laid down the foundation for understanding social interaction and culture. Although his book *Unified Theories of Cognition* was focused on the cognitive level, Newell speculated what implications his theory might have for the study of culture. Although these particular ideas did not receive much attention at the time, this paper argues that Newell's systems levels provide a valuable insight into the connection between brains and culture.

Keywords: Newell; Systems Levels; Culture; Knowledge Level; Cognitive Level

The concern over equity, diversity, and inclusion in cognitive science and other academic areas is linked to a corresponding rise in research seeking to understand the nature of discriminatory inclinations, both as they exist in the mind and in the culture. But what precisely does it mean to say that something exists both in the mind and in the culture? Moreover, what kind of constraints are placed upon an explanation that seeks to connect the contents of mind with the contents of culture, and how is this related to the physical brain?

This question is not new to psychology. In the landmark paper, "Social Psychology as History", Ken Gergan (1973) argued that the pervasive influence of culture makes it problematic to understand social psychology purely in terms of stable, natural laws. This is because the knowledge that defines the rules and the context for human interactions changes across situations, time and culture. It is important to absorb Gergan's full message; that the study of social interactions cannot be disentangled from the cultural context in which social interaction occurs. However, what is the relationship between culture and cognitive science? Cognitive science studies the cognitive level (information processing), which provides the bridge between the neural level (natural science) and the social level (history and culture). So the question can be reframed as, how does culture arise from neurons through cognition? Which leads to the question of neural realism. That is, why isn't everything reducible to neural activity and ultimately the evolutionary forces that shaped the brain? How is it possible to talk about culture and cultural forces without abandoning the commitment to ground our understanding of cognition in the physical brain? Fortunately, Allen Newell provided some theoretical scaffolding for tackling this problem. Here we argue that Newell's ideas on systems levels provide a useful way to situate cognition within culture as well as to understand the nature of culture itself.

Sui Generis Entities

In the cultural studies areas, such as anthropology, sociology, women's studies, critical race theory, queer studies, etc., there is an understanding of culture as *sui-generis*. Sui generis means, of its own kind. It is similar to the idea of strong emergence, or strong systems levels in Newell's scheme. In this sense, cultures, or cultural forces, exist and interact according to their own emergent rules. This can be illustrated by looking at the concept of racism. A racist individual can be understood as a person who endorses racist ideas. In this sense, racism exists at the level of individual beliefs or knowledge. However, systemic racism refers to a racist bias at the cultural level. In this case, a person with no racist beliefs can still perpetuate racism by participating in a culture that includes elements of systemic racism. However, this raises questions about the relationship between culture and the cognition of individual humans.

Newell's Systems Levels

It must be stated at the outset that while many aspects of Newell's overall framework are empirically grounded, the framework as a whole is on his account "speculative". It is a scientific argument, leveraging both what has been gleaned empirically in the experimental sciences and what has been discovered through engineering intelligent systems, all of which is aimed at the goal of achieving a "unified theory of cognition" that situates itself in a world where both brains and cultural forces can share equal claims to existence. As such, it is pioneering work, it is approximating work, and the foregoing is our interpretation of the framework itself and how it can be used to answer the question that everyone is trying to answer: how do we map the relationship between minds and culture?

For Newell, the human mind is engineered, via evolution, to have systems levels (Newell, 1994). Newell often exemplifies the properties of systems levels with reference to the design of computer systems. This is because computers have systems levels deliberately designed into them, so we

In J. Culbertson, A. Perfors, H. Rabagliati & V. Ramenzoni (Eds.), *Proceedings of the 44th Annual Conference of the Cognitive Science Society.*

know they exist. For example, most commonly, computers have a hardware level upon which an operating system level is engineered, upon which a software level (or app level) is engineered. Because it is engineered to have systems levels, software programmers need understand nothing about hardware engineering in order to do their job. We should also note that in neuroscience it is more common to talk about nested systems, but the idea is the same.

This can be compared to the human brain, where the neural structures are the hardware, the cognitive system is the operating system (i.e., a set of functional abilities realized in the hardware), and learned abilities are the software, realized through the cognitive system (note, we have updated this description somewhat and we also need to introduce the caveat that, in brains, learning can feedback to alter the cognitive and neural levels) .

In systems levels the components of one level are composed of or "realized by" the components in the level below. So while each layer expresses empirically different patterns of behaviors, they interface with one another effortlessly because a subset of interacting components at one level form a larger unified component at another level. Although this is a pretty abstract concept, the principle is intuitive when exemplified with real world examples. We know, for instance, that atoms interact with one another to build molecules, which interact with one another to build organelles, which interact with each other to build neurons, which interact with each other to build brains. None of this is disputed, and therefore the proposition that systems levels are part and parcel of human biology is not itself a speculative aspect of Newell's framework.

However, talk about levels quickly gets confusing. According to Wright and Bechtel (2007), "levels-talk is virtually threadbare from overuse yet [the] various conceptions of levels are rarely analyzed in any sustained, substantive detail" (p.55, see also Kersten, West, & Brook, 2015). One of the main sources of confusion is that we use the word "level" in two different ways: ontologies or epistemologies.

Ontologies are posited to exist in reality. They are made up of components, dynamics, and governing laws or principles that are presumed to physically exist. Empirically, quantum physics is widely regarded as the lowest level ontology, with other ontological levels built up from there. For example, chemistry is emergent on physics, biochemistry is emergent on chemistry, neurons are emergent on biochemistry, brains are emergent on neurons, and cognition is emergent on brains. Following on this, the claim that culture also exists as an emergent phenomenon, or sui generis entity, is an ontological claim. Note that in this type of system everything is built up from and grounded in physics.

Epistemologies, by contrast, are not levels in the world. In this way of thinking, epistemologies, or theoretical languages, are created because they are useful for understanding certain ontological phenomena. In this sense, epistemologies are in the mind of the researcher and not in the phenomena being studied.

To see the relationship between ontologies and the epistemologies used to understand them, let's take Lego as an example. We can use Lego to build different shapes and we can understand the shapes in terms of the types of Lego bricks used and how they were arranged. We could then take our shapes and combine them in specific ways to make new, more complex shapes. These new shapes can be considered as a new systems level because they were built on the old shapes and thus can be understood in terms of the old shapes, without reference to the original Lego blocks on which the old shapes were built. However, the language of analysis is still qualitatively the same, it is about shapes. Now imagine that we combine the shapes and make a Lego Turing Machine. We can still describe the Turing machine in terms of shapes, but it also exists sui generis, as a qualitatively new thing that requires a new type of language to describe it, information processing in this case.

This is a subtle distinction that Newell (1994) eventually cached out as systems levels and systems bands (see Figure 1). Systems levels can be understood in terms of being built on the components of the level below, without reference to how those components were built, whereas systems bands refer to adjacent groups of systems levels that are best described through a specific epistemology. For example, neural circuitry tends to be arranged in systems levels (or nested structures), with different functions, but the theoretical language used to describe all of them is the language of neuroscience. When a new language of understanding (epistemology) is needed, it signals that a new systems level has emerged that is unprecedented. This, according to Newell, marks the transition to a new band (see also Pylyshyn 1980, 1984).

TIME SCALE OF HUMAN ACTION			
Scale (seconds)	Time Units	System	World (theory)
10^{7}	months		SOCIAL BAND
10^{6}	weeks		
10^{5}	days		
10^{4}	hours	Task	RATIONAL BAND
10^{3}	10 min	Task	
10^{2}	minutes	Task	
10^{1}	10 sec	Unit Task	COGNITIVE BAND
10^{0}	1 sec	Operations	
10^{-1}	100 ms	Deliberate Act	
10^{-2}	10 ms	Neural Circuit	BIOLOGICAL BAND
10^{-3}	1 ms	Neuron	
10^{-4}	100 μs	Organelle	

Figure 1: Newell's levels and bands

However, the term, band, never caught on. Instead, people refer to bands (e.g., the cognitive band) as levels (e.g., the cognitive level). It is unavoidably awkward to introduce unused terminology, so we will sometimes refer to bands as levels (as in the cognitive level) and reserve the term band for situations where further clarity is needed. We will also refer to levels within bands as systems levels.

Another thing we need to consider before moving on is the issue of strong versus weak systems levels. The reason why computer system architectures are such a great exemplar of systems levels is because computers are deliberately engineered to embody a strong form of systems levels. The strongest form of systems level is one in which the interaction between its constituent components are all that need be considered in order to fully account for the behavior at that level. These levels are sealed off from lower levels in such a way that they can be called "state determined", which is to say that their behavior can be perfectly predicted knowing only the details of their current state (Newell, 1994). This perfect seal is also what allows for specialists to channel their vocation into one particular level and not have to worry about the dynamics of lower levels, unless of course, an error occurs that has knock-on effects that permeate upward. But errors are not the only reason why lower levels manifest changes in higher levels, for as Newell says, "there could be lots of ways in which phenomena from lower levels percolate upward, ways that are not to be described as errors" (Newell, 1994, pp. 118-119). Systems levels that by their very nature require the inclusion of components from lower levels to explain their behavior are said to be weaker systems levels. These are more typical of the systems levels that we find in nature. It is an open empirical question exactly how strong a systems level need be in order to usefully qualify as a systems level. It must demonstrate a pattern of novel behavior that is predictable to some significant degree, but exactly how much is not clear.

From Neurons to Culture

All of this has implications for our purposes here because it raises a very interesting empirical question: is culture an emergent level requiring its own unique theoretical language? Indeed, a growing number of scholars believe this to be true. Here we will map out culture by tracing through Newell's scheme, as illustrated in Figure 1.

The Neural Level

We will start our journey upward toward culture from the neural level. Here we take it that the existence of this level is uncontroversial.

The Cognitive Level

We would like to say that the cognitive level is also uncontroversial, but that is not exactly true as some people take a neural realism point of view and believe that only the neural level is "real." Here, Newell's claim that the cognitive level is real is an empirical claim. It is a claim that the neural architecture is arranged in systems levels and that these systems levels build to a point where the emergence of symbolic processing occurs, where representations encoded in neurons correspond to physical objects and events, allowing us to think about and simulate events in our minds by manipulating these representations. Here, we note that Newell had a very loose definition of symbols. In fact, Newell's only criterion for a symbol was distal access (Newell, 1994), which occurs when one brain module can send information to another brain module in such a form that it can be further processed at its destination. Therefore, symbols could be distributed, they could be fuzzy, or they could involve simulation. In other words, symbol processing is not limited to crisp propositions and formal modus ponens operations by this account. This is a claim that the cognitive level is an emergent phenomenon that is defined by the language of computational processing, but not a claim that we need to use a particular form of computational processing (Newell also had ideas about the form but that is a separate issue and should not be confused with the claim about systems levels).

The Knowledge Level

The most important thing to note is that Newell did not go straight from the cognitive level to social interaction and the level of social psychology. Instead, Newell had an intervening level that he variously called the knowledge level, the rational level, and the intendedly rational level. This level, which we will refer to as the knowledge level, is related to what Popper called the rationality principle and what Dennet called the intentional stance (see Young & West, 2018 for a review). The knowledge level can be understood in terms of agency and rationality. At the knowledge level, agents have goals and can be understood as rationally pursuing those goals, where rational is defined as bounded rationality (Simon, 1990). That is, the agent does its best with what it knows, it does not have perfect knowledge. Another important property of the knowledge level is that this is where semantic meaning occurs. For Newell, the cognitive level is only syntactic. Semantics, for Newell, exists in use, specifically it arises from pursuing goals. Thus, in Newell's scheme, meaning is always contextualized by the goals of the agent.

Finally, it is important to stress that by "goal," Newell was not referring to explicit, consciously held goals. What Newell meant by "goal" can be seen by examining SOAR, which is the cognitive architecture he created. As noted by Laird, Lebiere, and Rosenbloom (2017), SOAR, as well as other related cognitive architectures, uses working memory to represent the current state of the task. Production rules that match working memory contents are triggered and perform actions, including updating working memory. This creates a cycle that accomplishes the task (or solves the problem). The end goal is implicit in the chain of production rules that fire. This occurs at the cognitive level. However, it is only at the knowledge level that this is characterized as employing knowledge in pursuit of an end goal (unfortunately SOAR and other architectures tend to call the working memory

buffer that contains a representation of the current state of the task, the goal buffer, which is confusing). Intentionality, as Dennet would put it, or pursuing a goal, as Newell would put it, is an emergent property of the knowledge level.

In our opinion the knowledge level has been greatly misunderstood. Perhaps because its name, the "knowledge" level, people often think it is about how knowledge exists in the brain (i.e., how it is represented at the cognitive or neural levels). In fact, it is the opposite, the knowledge level was conceived of as a way of skipping over the problem of knowledge representation, by moving to a higher level. The claim is that knowledge exists in a form that can be detached from the way it is represented in individual agents. For example, the intelligent actions of a computer program can be described at the knowledge level (especially for expert systems) without reference to the specific software language they will eventually be coded in. As Newell (1994) notes, this ability to represent the same "knowledge" using different representational systems is behind our ability to transfer knowledge between agents (both humans and artificial intelligence) and to store it in artifacts (books, websites, etc.).

The Social Level

Newell also speculated about higher systems levels within the knowledge level band. Indeed, the ability to transfer knowledge between agents at the knowledge level would seem to support group cohesion and coordination for common goals. Newell speculated that group level processing could be limited due to the small bandwidth of communication between humans, compared to the bandwidth of communication within human brains (Newell, 1994). That is, Newell raised the issue that intelligence peaks in individual humans. In contrast, Ed Hutchins (Hutchins, 1995) has argued that humans rely heavily on distributed cognition for survival and to perform complex activities. For example, Hutchins (1995) gives a detailed example of a missile destroyer ship and how no individual on board has the full knowledge of how to sail the ship or complete the mission. In this case, the ship itself can be argued to be acting at the knowledge level. Distributed cognition is highly consistent with the knowledge level though, so it stands as a claim that the bandwidth problem can be solved and that relatively large groups of humans can function as intelligent entities. However, we also need to point out that groups are organized differently and it is likely that many groups would constitute only a weak systems level, relying partially or heavily on individual human knowledge level processing. That being said, it seems clear that the knowledge level can be applied to describing individual humans, groups, teams, businesses and institutions. In fact, it makes sense that humans would design their group intelligence systems based on their own individual intelligence.

The Cultural Level

Newell speculated that there could be higher level emergent bands, but he also speculated that the knowledge level may be the highest band, meaning that culture would be a systems level within the knowledge level.

> Why shouldn't the intendedly rational level just extend upward?…The issue of continued upward extension and perfection of rationality is whether groups can be described as knowledge-level systems — as a single body of knowledge with a set of goals — such that the group's behavior is explainable and predictable by its use of knowledge in the service of its goals (Newell, 1994, pp. 154).

In fact, both could be true, with some things we think of as culture residing within the knowledge level band and some things residing above the knowledge level band as a sui generis entity. Let us first consider culture as a phenomenon within the knowledge level.

Culture at the Knowledge Level

The knowledge level is characterized by the intention to use knowledge to move toward a goal. This can obviously be applied to groups, institutions, political parties, and movements with a stated goal. In this sense, cultural groups can be thought of as having goals and working toward them. However, an interesting question can be raised here. Since the individual humans within a group are at a lower systems level than the group itself, could the group have goals that the humans are unaware of? Certainly, it is the case that group leaders and influencers can exploit other group members by manipulating them to work on goals they are unaware of. Technically, though, if the systems levels were strong enough, the group itself could have goals that none of the humans involved are aware of.

If groups can have goals that some or all of their members are unaware of, then claims such as the patriarchy has the goal of oppressing women, or white supremacy has the goal of making the white race dominant, can make sense. It could also make sense to claim that individuals are supporting group goals, such as white supremacy or maintaining the patriarchy, by unwittingly participating in them. However, we need to raise two caveats here. The first is that it is not clear how humans, at a lower systems level in the knowledge band, can accurately ascertain the goals of groups at a higher systems level. The second is that, since humans use knowledge level thinking to predict other humans, we might fallaciously perceive that groups have intent or goals when it is not the case.

Culture above the Knowledge Level

It is important to remember that for culture to function as a sui generis agent, it must do so by co-opting the governing principle of agents at the knowledge level. This governing principle is known as the principle of rationality: "If an agent has knowledge that one of its actions will lead to one of its goals, then the agent will select that action" (Newell, 1982, pp. 102). Thus, a sui generis cultural level could be understood in terms of generating goals and knowledge,

which is different from using knowledge to pursue goals in the knowledge level.

Why do rational agents strive for the things they do? If you ask why we have the goals we do, we will always answer that question with reference to yet more goals. Why do you want a PhD? Because I want to become a professor. Why do you want to become a professor? Because I want to be filthy rich. Why do you want to be filthy rich? One can only refer to their goals with yet more goals, and once this line of questioning reaches its end, there is left a goal at the top that cannot be accounted for. This is known as the problem of the top goal.

It's not clear that human goal structures can be organized such that one goal emerges among all in the top spot. But at the very least, we appear to harbor a set of dominating goals that drive our behavior in various arenas of our life. When speculating about how culture as a level above the knowledge level may interface with rational agents in order to carry out its own agenda, perhaps it is through these top goals that it does so.

In addition to supplying us with top goals, another mechanism through which culture can control rational agents is by mediating the type and amount of knowledge endowed. It is one thing to have a goal, it is completely another to have a body of knowledge that enables the strategic pursuit of that goal. While it does appear that we are born into this world with some knowledge latently built-in and booting up in early development (e.g., object permanence), we acquire a great deal of knowledge through culture. Thus, we need to consider the mechanisms by which culture could create and disseminate goals and knowledge.

One mechanism is through aesthetics. In the same way that an aesthetic can drive an artist to pursue goals inspired by the aesthetic (the impressionists, for instance), it can also drive a larger scale culture by *resonating* with individual humans. When a person resonates with an aesthetic it can drive what kinds of vehicles they buy (Ram vs Prius), what kind of jackets they buy (leather vs hemp), and what kind of food they eat (steak vs tofu). Aesthetic is an interesting candidate because it clearly holds sway over rational agents, as seen by the successfulness of advertisement campaigns that leverage it, and it appears to be a high-level emergent phenomenon.

Another possibility is that culture operates by perpetuating grand narratives. Grand narratives are stories about what is best for the human race. Originally grand narratives existed in the form of myth and were further developed through religious thought. However, following postmodern philosophy, modern grand narratives are couched as political, economic, and ethical theories. For example, Marxism is a grand narrative, as is capitalism. Grand narratives are particularly interesting because they supply a top goal.

Also, we need to consider historical forces. Typically, what we think about when we hear the word "history" is sequences of events. But for thinkers like Hegel, Marx, and Fukuyama, history is a force with momentum and trajectory. These thinkers view the material world as being shaped through a dialectical process being carried out in the "realm of ideas" and "ideologies" (Fukuyama, 1989). For Hegel, the originator of this way of thinking, all of human behavior can trace itself back to prior ideas, and in this sense, is nothing more than the unfolding of the history of ideology. Taken from this perspective, the narratives, aesthetics, and values of culture are nothing more than the tip of the ideological arrow that has been sailing through time.

Finally, we also want to include evolution in this discussion. Newell speculated that Evolution was an even higher-level band than culture. The reason he placed evolution high in the scheme is the time scale. As Figure 1 shows, higher systems levels operate at a slower rate than lower systems levels, approximately by an order of magnitude. This is because the higher levels are inclusive of all the activities of all the lower levels. So, in the sense that evolution is occurring right now, it is occurring very slowly. Also, since Newell wrote this there has been considerable development in the field of cultural evolution. Memes (Dawkins, 1976) are the most well-known theory, but the field involves more than just memetic theory.

Concluding Remarks

As goal-directed agents, we are primed to view the world from the perspective of the knowledge level, or what Dennet called the intentional stance. Newell was aware of this problem, recognizing that "we humans are sitting within a level, looking upward, so to speak, at the social or historical system. We are components of the higher level and so we naturally see our own dynamics." (Newell, 1994, pp. 153). Due to this liability, we may tend towards ascribing intentions to sui genesis cultural entities when, in fact, goals do not drive behavior at that level. According to the analysis presented in this paper, if there are sui generis cultural forces, they create goals and knowledge but are not driven by them. That occurs at the knowledge level. However, importantly, some human groups, such as companies or political parties, do seem to exist within the knowledge level, with clear goals and knowledge sets. So it is important to be clear which level of analysis we are applying when we talk about groups. In our view, Newell's ideas provide a solid theoretical scaffold on which to integrate cognitive science with cultural studies.

References

Dawkins, R. (1976). *The selfish gene*. Oxford University Press.

Fukuyama, F. (1989). The end of history? *The National Interest,* (16), 3-18.

Gergen, K. J. (1973). Social psychology as history. *Journal of Personality and Social Psychology*, *26*(2), 309.

Hutchins, E. (1995). *Cognition in the wild*. MIT press.

Kersten, West, Brook (2015). Leveling the field: Talking levels in cognitive science. *Proceedings of the 38th Annual Cognitive Science Society Conference.*

Laird, J. E., Lebiere, C., & Rosenbloom, P. S. (2017). A standard model of the mind: Toward a common computational framework across artificial intelligence, cognitive science, neuroscience, and robotics. *Ai Magazine*, *38*(4), 13-26.

Newell, A. (1982). The knowledge level. *Artificial Intelligence*, *18*(1), 87-127.

Newell, A. (1994). *Unified theories of cognition*. Harvard University Press.

Pylyshyn, Z. (1980). Cognition and computation: Issues in the foundation of cognitive science. *Behavioral and Brain Sciences, 3*(1), 154-169.

Pylyshyn, Z. (1984). *Computation and cognition: Toward a foundation for cognitive science*. MIT Press.

Simon, H. A. (1990). Bounded rationality. In J. Eatwell, M. Milgate, P. Newman (Eds.), *Utility and probability* (pp. 15-18). Palgrave Macmillan London.

Wright, C., & Bechtel, W. (2007). Mechanisms and psychological explanation. In P. Thagard (Ed.), *Philosophy of psychology and cognitive science* (pp. 31-79). North Holland.

Young, J., & West, R. L. (2018) Exploring the reality of the knowledge level: Pragmatism embodied. *Proceedings of the 39th Annual Conference of the Cognitive Science Society*.