

REDUCTIONISM AND PRACTICALITY

Kevin Smith & Edward Vul

ABSTRACT: Like most domains of science, the study of the mind has been tackled at many scales of analysis, from the behavior of large groups of people (economics and ecology), to the diffusion of ions across cellular membranes (molecular biology and biophysics). At each of these scales, researchers often believe that the critical phenomena of interest, and the most powerful explanatory constructs and mechanisms, reside at their scale of analysis, with finer scales argued to be incapable of predicting the interesting phenomena, while coarser scales are purported to miss critical mechanistic subtleties. Here we argue by analogy that, for better or worse, researchers at all scales are correct: phenomena at each scale of analysis are intractable from other scales; thus, while reductionism is a useful scientific goal, it will not obviate the need for macroscopic research, constructs, and formalisms.

KEYWORDS: reductionism; levels of analysis; philosophy of science

ALL MODELS ARE WRONG, BUT SOME ARE USEFUL – *George Box*

At what scale should we study human behavior and the brain processes responsible? Fields at adjacent scales of analysis find themselves in similar disagreements, with every field condemning more abstract scales of analysis for failing to describe “how” various processes work, and more minute scales for missing the forest for the trees. Conveniently, the purportedly appropriate scale of analysis for studying “mechanisms” of mind and behavior usually coincides with that of the speakers’ own research program.

Computational cognitive scientists argue that understanding the brain requires a more abstract description of the organisms goals and available information, and claim cognitive psychologists face an insurmountably under constrained task in trying to describe the function of the whole organism by cobbling together isolated algorithms. Cognitive psychologists in turn condemn computational cognitive scientists for working with idealized learning models and glossing over specific errors and deviations from these idealizations found while studying particular psychological mechanisms.

Cognitive psychologists assert that constraining processes and algorithms from neurological systems discovered through imaging data would be impossibly complex, while cognitive neuroscientists point out that the constructs that cognitive psychologists study cannot identify the neural mechanisms that necessarily underlie those processes.

Cognitive neuroscientists themselves argue that characterizing the aggregate function of a brain area from the behavior of individual cells is hopelessly difficult, while electrophysiologists contend that cognitive neuroscientists merely show which areas are active, and cannot characterize the circuitry and mechanisms within those areas.

Electrophysiologists in turn believe that studying the function of a cell based on its connections with its neighbors would be problematic to accomplish, while connectomicists argue that the connections among neurons are the mechanisms of brain function.

Connectomicists assert that studying neural connections in terms of the molecular processes of synaptic formation is irreducibly complex, while cellular and molecular neuroscientists contend that understanding these connections requires studying the formation and maintenance of synapses at the molecular level.

There is a common thread throughout all of the across-level disputes: the higher scale of analysis (greater abstraction) maintains that extracting their phenomena of interest from lower scales of analysis is impractical and hopeless, while the lower scales (less abstraction) insist that the constructs used as explanatory mechanisms at the higher scales simply do not exist when scrutinized, and conceal much more complicated and varied, smaller scale phenomena.

Who is right? Here we contend that both sides of these many arguments are correct, and that reductionism, while a useful scientific goal, will not eliminate the need for the more abstract scales of analysis.

HOW?

What does it mean to answer a “how” question? *How* do people learn? *How* do we speak? *How* do we emote? Locomote? See? Researchers in many subfields of behavioral, cognitive, and neural sciences, all purport to be studying the “mechanisms” of behavior that answer these how questions. Unfortunately, one scientist’s mechanism, is merely a *phenomenon* in need of a mechanism to another. Thus, while one researcher may consider their explanations to be satisfying answers to how questions, to others that explanation itself needs to be explained to answer “how”.

Not only do the basic units and phenomena differ across scales, but these scales are also conducive to different formalisms: when describing the behavior of groups of

individuals, microeconomics offers elegant equations for the interactions of supply, demand, and price, but cannot describe how a single individual will react to economic changes. When studying the behavior of individuals, formalisms at Marr's computational level (e.g., Bayesian statistics) can specify our goals, available information, and prior knowledge we use make inferences and choose actions, but then we are agnostic to the process the mind uses to carry out this information processing. When studying the microstructure of behavior, such as variations in response times, constraints on memory, and processing speed, we must describe human cognition at the algorithmic level using the language of computer science, describing how people represent data, and what procedures operate over these representations to make the required computations. If we focus on the scale of neuronal ensembles, we would instead adopt the language of electrical engineering, and talk about the neural signals, systems, and circuits that are the physiological instantiations of the algorithmic description. We could reduce further to the level of biochemistry, where we describe the individual neurotransmitters, ion channels, and chemical gradients that allow neurons to pass information between one another and generate action potentials. Of course, we needn't stop there, since those individual neurotransmitter molecules and ions are comprised of atoms and subatomic particles.

So, how do we decide at which scale of analysis, or level of abstraction to operate? Before offering an answer, let's consider an engineered system, where we (well, maybe not us specifically – but someone) can characterize the system at all relevant scales.

SCALES OF A SMARTPHONE

“How” does my smartphone correct my typos? As in the case for the how questions about human behavior, we can answer this at many scales of analysis, and the appropriate scale depends crucially on the question we are asking.

What will my typo be replaced with? At the information scale of analysis, the phone has access to a dictionary of acceptable words, and has some distance metric between two strings (a distance metric based on likely data entry errors made on a qwerty keyboard). Based on this distance metric, the phone can identify which dictionary words are closest to the entered string, and thus suggest replacements (perhaps weighting the distance metric by the probability of a given word within the context of other words). The dictionary and distance metric determine which words will be the top suggestions for a given string.

How fast will it make the correction? At the algorithmic scale of analysis, the phone has particular search algorithms that it uses to pick out candidate subsets of words from the dictionary, and to compute the distance between a string and those candidate

words. These algorithms (instantiated in machine code) determine how much memory and relative CPU time it will take to produce candidate corrections for various strings.

What happens if the phone breaks? At the scale of circuitry, we may specify the locations and arrangement of micro transistors in the phone, its flash memory store, and where on that array the dictionary is stored, the architecture of the CPU, and how it is connected to the display, etc. At this level of description, we might be able to indicate what parts of the spell-correction process would break if we were to snip some circuitry within the phone.

How much energy does this process use? At the scale of materials, we might specify the conductive properties of the circuitry connecting memory stores and comprising the CPU architecture, as well as the chemical properties of the battery providing the energy required for the system to operate at all. At this scale of analysis, we might be able to predict why spell-checking drains more battery power when the weather is cold, or why the phone gets warm when we ask it to do lots of spell-checking at once.

Each of these questions is best suited to a different level of analysis. It would be wildly impractical to reason out the abstract distance metric and contents of the dictionary from machine code, and determining it from bare circuitry, while theoretically possible, would be hopeless in practice. Conversely, an understanding the distance metric and algorithm simply offers no language to express material or circuit properties that could describe phenomena at that scale.

In short, each scale of analysis, or level of abstraction, is useful – the more abstract scales allow us to predict higher order phenomena without simulating intractably complicated systems; the less abstract levels of description allow us to capture phenomena that may be inexpressible at greater abstractions. We believe that restricting the study of the brain to a single “correct” level of analysis would be akin to restricting a smartphone designer to hiring only materials scientists or electrical engineers or software developers – each provides solutions to different sets of problems, and a smartphone could not be fully understood, much less developed, without input at all levels of analysis.

SCALES OF A MEMORY SYSTEM

“How” do we remember things? Unlike the smartphone example, we do not know how the brain works, so we do not have a roadmap that segments levels of analysis and defines how they connect. Nonetheless, different questions lend themselves more readily to different levels of analysis, as demonstrated by successes in memory research:

What will we recall if cued with a particular item? Just like studying a spell-check process, this question requires us to know the contents of memory (the ‘dictionary’) and the process by which items are queried. Researchers have been investigating how the query process works, including how combinations of cues and context might influence recall (Polyn, Norman, & Kahana, 2009) or how we search through our stored memories to find close associations to a cue (Abbott, Austerweil, & Griffiths, 2012). These theories are intentionally agnostic to the biological bases of memory: knowing where in the brain memories are stored or the biochemical reactions that lead to neural changes will not describe the retrieval process adequately.

How should we provide information in a classroom to increase learning? If we want to apply psychological research to education, then we must focus on the process by which new information is stored in memory and old information is forgotten. Here researchers investigate the optimal timing and method of presenting and testing material in order to ensure students learn and retain information (Roher & Pashler, 2010). Knowing how the mind stores and queries information can constrain these processes, but even a perfect knowledge of this cannot predict the time profile of adding information to these stores. Similarly, this learning is achieved through a complex process that takes perceptual information, processes it, and stores it; knowing the biological bases of this process would in theory allow us to replicate this time course, but would require significantly more computation to provide us with the same answer.

What brain areas should be spared during surgery to avoid memory deficits? If we want to know where in the brain memories are processed and stored, our cognitive theories of the storage and retrieval of memory are uninformative – instead we must study the biological bases and memory circuits involved in this processing. For instance, we may rely on decades of research from patient studies (Scoville & Milner, 1957) and brain activity patterns (Smith, Wixted, & Squire, 2011) that have suggested that the hippocampus is required for forming new memories.

What drugs might promote better memory? At this level, even knowing the brain areas and circuits underlying memory will provide little help. Instead, we need to understand the biochemical bases beneath these brain regions: what pharmacological agents, introduced at what time, might increase long-term potentiation (Stella, Schweitzer, & Piomelli, 1997), inhibit synaptic degradation (Pak & Sheng, 2003), and otherwise modulate the synaptic changes that underlie learning and memory (Davis & Squire, 1984)?

TRADEOFFS IN SCALES OF ANALYSIS, AND REDUCTIONISM

Even if we consider physics—a model of reductionist success—we find that although we can reduce macroscopic constructs to their microscopic mechanisms, all the way down to subatomic particles, we do not throw out the formalisms at the more abstract levels. When dealing with a higher order abstraction (like classical mechanics), we will fail to account for some subtleties that would be captured at a finer scale (quantum interactions), which could end up playing an important role in the phenomenon of interest. When dealing with lower abstractions (such as particle physics), we face a vast computational challenge when trying to describe higher-order phenomena (like how a ball will bounce). Thus, physical models at different levels of abstraction have proved to be more or less useful depending on the phenomenon of interest, so different abstractions are emphasized in astrophysics, mechanical engineering, electrical engineering, and quantum computing. Climate and meteorology models do not attempt to predict weather from individual quantum particles, because it would be a futile computational endeavor; they do not even consider the behavior of individual particles, or the statistical mechanics of groups of particles – instead they adopt volumetric analyses of the temperature and pressure of kilometer, or hundreds-of-kilometers scale regions of the atmosphere.

Yet despite the large numbers of scales and models used in modern physics, reductionism has been a necessary part of the research enterprise – higher level models can place constraints on lower levels, and lower level models can provide additional insight into higher level. Quantum mechanics, for instance, would have never been accepted as a theory had it not scaled up to classical mechanics on a macroscopic level. Conversely, findings from quantum mechanics have led us to understand superconductivity, which has lent itself to many novel uses in electrical engineering.

Reductionism must hold a similar place in the study of the brain. We are reassured of our scientific models at higher levels of abstraction (like trichromacy – the theory that human color vision is three dimensional) when those models may be derived from properties at lower orders of abstraction (the existence of three cone types). Similarly, we are reassured that we are measuring relevant properties of complexly interacting elements (like receptive field size of V1 cells) when those properties can be simplified to abstractions about the important behaviors of the system as a whole (cortical magnification and the falloff of acuity with eccentricity). Thus, connecting two scales of analysis validates models at both levels of abstraction, so there is a scientific demand for a single unified model of human behavior, cognition, and neuroscience by reducing cognitive theories to their biological underpinnings. However, even when the levels of abstraction are united, such that theoretical constructs at higher levels of abstractions may be reduced to their

mechanisms at finer scales throughout the full hierarchy of analysis, we should not expect that one level of description will emerge as the most fundamental, useful, or practical.

CLOSING

When studying human behavior, we face a salient problem: we really have very little idea of how it all works, at any scale. The models and theories we work with at various scales of analysis are our best current approximations – subject to infinite revision and refinement, happening weekly. So, at this point it may not be obvious what a final understanding of human behavior will look like. Self-interested researchers often anticipate that the final, fundamental description of the “mechanisms” of behavior will be cast at their level of analysis. However, if we are to take any lessons from physics and engineering, we should note that is not what happens in these fields that have succeeded in characterizing their domains of interest – these disciplines yield multiple characterizations at multiple scales of analysis, with the abstract constructs at higher levels of description reducing to the interaction of less abstract constructs at lower levels of description. Moreover, even once such a chain of reductionism has been completed throughout all scales, we do not see that one level of abstraction dominates in its utility, or practical application: we aim to predict and manipulate phenomena at many different scales, and to do so, we must use an appropriate level of abstraction for each. Using too microscopic a scale of analysis yields intractable computational problems, and using too macroscopic a scale of analysis misses critical details of the phenomena of interest. We should expect the same reliance on all scales of analysis when we have figured out how to predict and manipulate human behavior.

Department of Psychology
University of California
San Diego

Address correspondence to evul@ucsd.edu

BIBLIOGRAPHY

- Abbott, J. T., Austerweil, J. L., & Griffiths, T. L. (2012). Human memory search as a random walk in a semantic network [Journal Article]. *Advances in Neural Information Processing Systems*, 3050-3058.
- Davis, H. P., & Squire, L. R. (1984). Protein synthesis and memory: A review [Journal Article]. *Psychological Bulletin*, 96(3), 518-559.

- Pak, D. T., & Sheng, M. (2003). Targeted protein degradation and synapse remodeling by an inducible protein kinase [Journal Article]. *Science*, 302(5649), 1368-1373.
- Polyn, S. M., Norman, K. A., & Kahana, M. J. (2009). A context maintenance and retrieval model of organizational processes in free recall. [Journal Article]. *Psychological Review*, 116(1), 129-156.
- Roher, D., & Pashler, H. (2010). Recent research on human learning challenges conventional instruction strategies [Journal Article]. *Educational Researcher*, 39(5), 406-412.
- Scoville, W. B., & Milner, B. (1957). Loss of recent memory after bilateral hippocampal lesions [Journal Article]. *Journal of Neurology, Neurosurgery, and Psychiatry*, 20(1), 11-21.
- Smith, C. N., Wixted, J. T., & Squire, L. R. (2011). The hippocampus supports both recollection and familiarity when memories are strong [Journal Article]. *The Journal of Neuroscience*, 31(6644), 15693-15702.
- Stella, N., Schweitzer, P., & Piomelli, D. (1997). A second endogenous cannabinoid that modulates long-term potentiation [Journal Article]. *Nature*, 388(44), 773-778.