Finding the Systematicity of Language in the Structure of Perceptual Experience

Alison –
http://sheesania.com/

17 May 2017

Consider these words you are reading, arbitrary collections of sound whose connection to the actual experiences they evoke is tenuous at best. Yet working from that huge, interconnected jumble of experiences juxtaposed with bits of language that you have had over many years of life, you have been able to learn distinct words, rules of grammar, and systems for effectively communicating with language in endless novel situations. Trying to reproduce this feat of richness and flexibility, many artificial intelligence researchers today are considering neural networks, modeled after the human mind. But neural networks struggle to generalize grammatical rules and structure out of the webs of linguistic information they are trained with. I propose that artificial neural networks can face this challenge and better learn and process the systematic grammar of natural languages by drawing on the rich structure of the real world in perceptual information, because the structure we discover in our perceptual experiences is what we use during our human language learning process to build up systematic concepts around the new words we encounter.

I will begin by introducing the structure of artificial neural networks (ANNs) and the connectionist philosophy of mind underlying them, and how researchers use ANNs with computers for natural language processing (NLP). I will also outline the systematicity challenge to connectionism and its implications for NLP with neural networks. Then I will argue that some of the structure necessary to systematicity comes from structure in the physical world, not just structure in the architecture of the mind. Following from this, I will reason that a neural network can become more systematic and thus process natural language better when it draws on information from the real world to access this structure.

Background: Connectionism and the Systematicity Challenge

Connectionism

Connectionism is a theory of cognition that tries to model the mind and thought through artificial neural networks, interconnected webs of individual nodes based loosely off the biology of the brain. These nodes are divided into three types of layers: an input layer, one or more hidden layers, and an output layer. In a human brain, by analogy, the input layer might be sensory neurons, the output layer motor control neurons, and the hidden layers the rest of the neurons that contribute to the thought process between sensation and action. Each neuron can receive, process, and send activation values. Thought begins when input units receive initial activation values, such as sensation. Each unit then sends a signal with its activation value to the units in the next hidden layer it is connected to, and then each of those units processes the combined signal it is receiving from its various connections. Every connection between units has a weight making the activation values sent along it stronger or weaker. The receiving unit simply adds the activation values it has received together, adjusting for the weights, and then sends that new activation value along to the other units it is connected to in the next hidden layer. Ultimately, after traveling through the input and hidden layers and being shaped by the different connection weights, the signal reaches the output layer. The pattern of activated nodes in the output layer indicates a particular result, like a movement. ANNs learn the right activation patterns by adjusting connections between units, whether automatically (unsupervised learning) or through manual adjustment by a human (supervised learning).[1]

This connectionist model stands in contrast to the classical model of cognition and mental representation, which sees the mind as a computer that processes distinct symbols using fixed, regular rules. Connectionism, meanwhile, does not represent concepts through distinct, localized units like classical symbols, but by particular distributed patterns of individually meaningless units activating all over the output layer. Instead of using a series of regular rules to handle input, neural networks have numerous connections of different weights between units that all interact together.[2]

This method of representation and processing give neural networks some unique strengths

[1] James Garson, The Stanford Encyclopedia of Philosophy, Winter 2016 ed., s.v. "Connectionism." Stanford: Center for the Study of Language and Information at Stanford University, 2016. https://plato.stanford.edu/archives/win2016/entries/connectionism/ (accessed February 25, 2017).
[2] Ibid.

and weaknesses relevant to NLP. For one, they can cope with corrupted or noisy input and make reasonable guesses from it like human cognition does, because as long as the input produces roughly the right activation pattern, the net can still capture the correct meaning. This makes them well suited for speech recognition, where the input sounds are never perfectly clear, and for handling exceptions in grammar like English's irregular past tense verbs. Connectionist representation can also efficiently carry information about the relationships and similarities between concepts, because related concepts often have similar activation patterns. This can give neural nets an edge in recognizing semantic relationships between words. However, while neural networks excel at matching patterns, they struggle more with following rules. They can capture similarities between instances with similar activation patterns, but they do not lend themselves to recognizing the underlying general, regular rules and structures that apply across many different instances.[3] This becomes an issue when dealing with all the regularities in grammar and phonology.

The Systematicity Challenge

Recognizing the weakness of neural networks in dealing with regularity, Fodor and Pylyshyn describe a concept called "systematicity" and show how language, and by extension cognition in general, is pervasively systematic. They argue that connectionist models are not intrinsically systematic, and thus aim to prove that connectionism is not a viable model of the mind.[4] But given their demonstration of systematicity in language, neural networks not being naturally systematic would also mean that they face a fundamental structural flaw when trying to process language.

According to Fodor and Pylyshyn, "systematicity" means that "the ability to produce/understand some sentences is intrinsically connected to the ability to produce/understand

[3] Ibid.
[4] Jerry A. Fodor and Zenon W. Pylyshyn, "Connectionism and cognitive architecture: A critical analysis," *Cognition* 28, no. 1 (1988): 3-71.

certain others."[5] They contrast memorizing sentences out of a phrase book to actually learning the language, where the speaker is able to form and understand novel sentences by combining known words and syntactic structures. If speakers understand "John loves the girl," then they must also understand "the girl loves John." They understand the syntactic formation (noun phrase, verb, noun phrase) and the semantic pieces ("John", "love", "girl"), and so they can *systematically* derive new sentences using those pieces, like "the girl loves John."[6]

The classical model of cognition would represent sentences like this in a structured fashion, with hierarchical trees of syntax elements holding individually meaningful semantic pieces. Thus, sentences with the same structures or semantic elements are immediately apparent, and it is easy to systematically form new sentences by swapping out elements in the tree.[7] In contrast, connectionist representations are simply piles of "microfeatures"[8] indicated by units. Without any structure or distinct subparts, "John loves the girl" and "the girl loves John" have no particular relation. They are both unique blobs of activation patterns, not trees with the same shapes and elements. Therefore, even if a neural network can process "John loves the girl," it cannot necessarily process "the girl loves John."[9] It lacks the intrinsic systematicity of a classical model.

Fodor and Pylyshyn agree that it may be possible to implement a connectionist network that properly recognizes all representations with the same structure.[10] But it would retain a fundamental disadvantage, because it naturally treats language as the ability to use some number of independent sentences instead of the ability to *generate* sentences using a store of words and syntax rules, like in a classical model. Neural networks will always struggle with systematicity because it goes against their own grain, and thus, given the systematicity of language, they will always struggle with NLP.

---

[5] Ibid., 37.
[6] Ibid., 37-39.
[7] Ibid., 38-39.
[8] Ibid., 19-20.
[9] Ibid., 38-40.
[10] Ibid., 64-66.

Fodor and Pylyshyn's prediction has been borne out in several experiments attempting to design neural nets for NLP. Pinker and Prince, for instance, examine a neural net constructed by Rumelhart and McClelland meant to produce past tense versions of English verbs.[11] Their net generalizes irregular inflections far too often, failing to recognize them properly as exceptions to a general rule.[12] Hadley looks at six different nets developed by different researchers, all of which can only recognize a word in a syntactic position it has been seen in before, at best.[13] This is a far cry from human language learning, where we can easily use a new word in almost any syntactic formation after only hearing it once! The net Hadley considers to be closest to being fully systematic also depends on existing structure, since instead of natural sentences, it uses neatly parsed representations of them.[14] It "thus assumes the existence of an external agent who has already discovered and encoded the compositional structure of a set of representations."[15] If even quasi-systematicity requires "an external agent" that Hadley finds implausible,[16] then it seems unlikely that connectionist networks could reach strong systematicity and effective language learning.

Overall, Fodor and Pylyshyn's work and the experiments responding to it seem to indicate that systematicity is a real challenge for NLP with neural networks. Fodor and Pylyshyn show how systematicity is a key component of human language, but does not come naturally to connectionism. Neural networks have, in fact, had many practical problems with systematicity as they try to recognize grammatical rules and use new vocabulary. But if they could somehow overcome this fundamental flaw, their unique capability for pattern matching and flexibility could give them an advantage in NLP over classical architectures.

---

11 Steven Pinker and Alan Prince, "On Language and Connectionism: Analysis of a Parallel Distributed Processing Model of Language Acquisition," *Cognition* 28, no. 1 (1988): 73–193. http://users.ecs.soton.ac.uk/harnad/Papers/Py104/pinker.conn.html (accessed April 27, 2017).
12 Ibid.
13 Ibid., 254-270.
14 Ibid., 266.
15 Ibid., 270.
16 Ibid.

Argument: Perceptual Structure and Systematicity

According to Fodor and Pylyshyn as well as Hadley, the main requirement for systematicity is for a model to recognize abstract structure and categories such as grammatical formations and word categories, not just individual instances and their associations. In a classical model, this structure is present because its method of representation, hierarchical arrangements of symbols, requires it. However, Hadley's comment about "an external agent" suggests that there could be other sources of structure besides the nature of the model. Pinker and Prince's critique notes that "Rumelhart and McClelland made up for the model's lack of proper rule-motivated structure by putting it into a teaching environment that was unrealistically tailored to produce much of the behavior they wanted to see. In the absence of macro-organization the environment must bear a very heavy burden."[17] Could structure in the environment, not the model itself, somehow provide the abstract structure and categories necessary for systematicity?

Language as Labeling

As they consider the cognitive processes underlying how children first learn language, Dennett and Clark offer some insight into how the environment may play a role in learning systematicity. They portray child language acquisition as, in essence, the child learning how to label and manipulate knowledge he or she already has regarding the world, figuring out how to use language to structure and categorize existing information.[18] This implies that the structure of language is coupled to details of the physical world.

Dennett argues that the minds of animals and prelingual infants behave like nonsystematic connectionist networks. Such nonlinguistic creatures have extensive knowledge and are good at

[17] Pinker and Prince.
[18] Daniel C. Dennett, "Learning and labeling," *Mind & Language* 8, no. 4 (1993): 540-548. https://ase.tufts.edu/cogstud/dennett/papers/learn&la.htm (accessed April 2, 2017); Andy Clark, "Material Symbols," in *Supersizing the Mind* (New York: Oxford University Press, 2008), 44-60.

recognizing familiar patterns, but they often cannot generalize to unfamiliar data or isolate specific details. They cannot recall and reflect on their own knowledge, and so it is "opaque," not self-consciously "wieldable."[19] "The 'wisdom' of [a cuckoo's] behavior is in some sense embedded in the innate wiring…but the cuckoo hasn't a clue about this rationale."[20] It is trapped in "interwoven connectionist nests"[21] without structured symbols and thus without systematicity, making abstract higher-level thinking impossible. Clark describes an experiment in which researchers showed babies toys hidden in a room and then had the children look for them. The toys' locations were marked by a unique combination of the color and shape of a wall, but the prelingual babies only used wall shape to look for the toys, seemingly unable to combine two different kinds of cues even though they could recognize color and shape independently.[22] This sounds much like a spatial version of a lack of systematicity. The infants have knowledge, but cannot freely manipulate it.

According to Dennett's picture of language acquisition, children then stumble upon linguistic utterances, initially as "found objects"[23] much like any other perceptions they experience. The child repeats them and associates them with particular contexts, slowly becoming familiar with them.[24] But words are special because they are labels: distinctive, memorable, and arbitrary, each word "wear[ing] the deliberateness of its creation on its sleeve"[25] to make its purpose clear - it is a tool to reduce the world's complexity. As words become attached to known perceptions, their distinctiveness makes them "ready *enhancers* of sought-for associations that are already to some extent laid down in the system."[26] Children start to figure out how to use a label to pull a set of associations from their memories. Soon the process is so seamless that by manipulating words in a

---

[19] Dennett.
[20] Ibid.
[21] Ibid.
[22] Clark, 48-49.
[23] Dennett.
[24] Ibid.
[25] Ibid.
[26] Ibid.

language, they can also cognitively manipulate the concepts those words extract out of their experience.[27] Now with words to make structure and categories out of their "interwoven connectionist nests," systematic thought is possible.

Clark delves more into how words function as labels, "cognitive scaffolding"[28] that provides structure to and thus control over perception. A label lets us focus on particular relevant perceptions, shaping our experiences into different structures.[29] It also pulls those perceptions out of their original associational contexts so we can manipulate them freely, combining them with other concepts.[30] Language can help us anchor and coordinate our constantly shifting perceptions. A connectionist mind, like a child's, is highly context sensitive and elastic, but a tool like language can stabilize it[31] and confer "the ability to reliably follow trajectories in representational space and to reliably lead others through certain trajectories."[32] Language is so powerful because it draws structure and distinct symbols out of our experience of the world.

Imposing Structure or Discovering Structure?

Dennett and Clark suggest that our basic experience of the world is a stream of mental states, perceptions embedded in a context of associated perceptions. They conclude that we gain more control over these experiences by imposing labels and structure on them with language, organizing them so we can turn our pattern-matching machinery onto higher-level problems. To them, the meaning of language is "self-engineered,"[33] whatever helps us survive in our environment. Consequently, learning language is not actually discovering true structure in the world, but creating one's own structure and imposing it on the world. By their logic, language does not involve any

27 Ibid.
28 Clark, 44.
29 Ibid., 45-50.
30 Ibid., 50-53.
31 Ibid., 53-54.
32 Ibid., 53.
33 Ibid., 60.

normative process that discovers truth about the nature of the world. It is just the brain engineering perceptions in any way that is conducive to our cognitive processes and thus our evolutionary survival, and so language does not necessarily reflect truth. If this is the case and the real structure of the world is not actually the source of structure in language, then as we develop a neural network for NLP, we are left with the same problem of how to make structure out of connectionist associations. Must we try to haphazardly generate structures that are somehow "useful" in an evolutionary sense?

Dennett and Clark's biological, evolutionary understanding of meaning in language flows from a broader naturalistic view of reason that Dickerson explores.[34] Reason, like language, considers the world and appears to derive something from it: a proposition in this case rather than a label. But Dickerson points out that in a strictly physical conception of the mind, the brain is merely following its biological programming when it uses reason. Reason is not a normative process discerning truth about the world, either. Once again, it is just the mind engineering perceptions according to particular processes that have happened to be helpful for evolutionary survival.[35] The propositions that reason results in are created by the mind, just as the structure that language provides is created by the mind, and thus neither are based on any kind of outside truth.

But this stands in stark contrast to a Christian view, in which the world is not entirely physical but also has a spiritual realm. Humans are both spiritual and physical beings, and so our reason is partly spiritual and can function beyond the closed causal system, escaping our determined biological programming.[36] More significantly, we were given this reason by God, the perfect "divine Reasoner."[37] In a "miraculous"[38] fashion, He supports our reason with His own perfect ability to discern truth.[39] Free from a physical chain of causes, divine revelation assisting us, and made in the

[34] Matthew Dickerson, *The Mind and the Machine* (Grand Rapids: Brazos Press, 2011), 91-117.
[35] Ibid., 98-106.
[36] Ibid., 159-160.
[37] Ibid., 160.
[38] Ibid.
[39] Ibid., 160-162.

image of God the Reasoner ourselves, we can access truth.[40] Our reason is still fallen and imperfect because of sin,[41] but nevertheless we have a normative tool for reaching truth that often succeeds.

So according to a Christian perspective, correctly applied reason does not create propositions based on what is best for survival, but *discovers* propositions from some normative truth. Labeling is a kind of reason that derives propositions specifically about the structure and relations of experiences, and so by extension, a Christian view of labeling as a subset of reason would be that it discovers structure rather than creating it. Again, this indicates that there is some normative structural truth out there that we can access. So if children learn language by associating experiences with particular symbols, they are actually detecting some true cohesiveness and conceptual unity in those experiences - not just putting a set of unrelated perceptions together because it is cognitively convenient to do so.

Discovering Systematicity using Perceptual Information

According to a Christian theory of labeling, then, acquiring language is learning how to label experiences of the world, finding structures and categories God established in the world to build up linguistic structures and categories. The development of systematicity is thus dependent on our perceptual knowledge of the world that contains the systematic structures. So I suggest that NLP neural networks struggle with systematicity because as they attempt to grasp the structure and categories of language, they do not have that store of perceptual knowledge that is the source of the structure and categories. Connectionist pattern-matching cannot develop systematicity because the systematic patterns are not there to recognize. If, however, an NLP neural network had access to associated perceptual information along with language samples, I hypothesize that it would be able to recognize structure and categories much more effectively and thus behave more systematically.

[40] Ibid., 159-162.
[41] Ibid., 162-165.

Jansen and Watter present just such an experiment in their effort to model how an infant acquires language per the semantic bootstrapping theory, which posits that infants use an extensive conceptual knowledge of the world to rapidly learn new linguistic symbols.[42] Jansen and Watter designed an artificial neural network and had it learn unsupervised from "grounded" sentences formed of words marked with sensorimotor features. These features imitate existing perceptual concepts infants could plausibly recognize as they hear words, such as "is red" or "decreases hunger." To train the network to build associations and recognize these words, they used a function for finding semantic similarity based on how many features two representations have in common.[43]

After training the network, they looked at its underlying association structure for evidence of systematic word categories and syntactic formations.[44] Then they measured how well it could process the grammar of some new sentences with entirely unfamiliar words, though the words are still marked with sensorimotor features. As a control, they also trained an identically-structured network with ungrounded input and tested it against similarly novel ungrounded input.[45] The results are very positive. When learning words, the network developed similar activation patterns of members of grammatical categories like nouns or verbs, purely based on sensorimotor features.[46] Underlying activation patterns for sentences were also very similar for trained and novel sentences with the same grammatical structure, even when they shared no words. In the test of grammar processing, the grounded network did much better than the ungrounded according to the standard GPE measure of ungrammatical predictions, 1.9% mean GPE for the grounded network versus 43.1% GPE for the ungrounded one.[47] This performance was equal to or better than the best NLP

[42] Peter Jansen and Scott Watter, "Strong systematicity through sensorimotor conceptual grounding: an unsupervised, developmental approach to connectionist sentence processing," *Connection Science* 24, no. 1 (March 2012): 33.
[43] Ibid., 33-35.
[44] Ibid., 35-37.
[45] Ibid., 37-42.
[46] Ibid., 36-37.
[47] Ibid., 42-45.

networks at the time the paper was written.[48]

Jansen and Watter acknowledge that their input data is oversimplified, and this kind of sensorimotor grounding may not be as relevant to later linguistic development involving more abstract categories and grammar.[49] But overall, their model's success in recognizing distinct categories of interchangeable words and syntactic structures leads them to "posit that systematicity is not a property of the cognitive system alone, but rather...it can be a property of the pairing of an appropriately sensitive computational/representational system with an informative input set."[50] This conclusion and the rest of the experiment supports my answer to the systematicity challenge for connectionist NLP. Labeling theory says that children learn systematic structure and categories by interacting with perceptual experiences, which Jansen and Watter's neural net does as it figures out grammatical categories like nouns and verbs based on sensorimotor features designed to mimic perceptual experiences. My Christian variation of labeling specifies that children discover structure from these experiences, rather than creating it. Since Jansen and Watter's training algorithm was unsupervised, the net developed the correct grammatical structures with little direction on what kind of structure to form, so it is likely that it actually discovered that structure in the data and did not just generate it. Overall, the net's interaction with perceptual information significantly boosted its ability to learn systematic structure and behave systematically.

Conclusion

NLP neural networks often flounder when trying to handle systematic rules of grammar, since they are not naturally sensitive to structure. Fodor and Pylyshyn present this as a larger issue for connectionism as a model of the human mind. Explaining human language learning, Dennett and Clark argue that the mind stumbles on labels in the form of words, and then figures out how to

[48] Ibid., 45-49.
[49] Ibid., 49-50.
[50] Ibid., 48.

use them to gather and sort perceptual experiences into systematically manipulable symbols. I argue that in a Christian rather than naturalist view of labeling, the mind uses labels to discover structure in its experiences and build up symbols. This theory implies that given the right perceptual context along with symbols, an ANN could detect the structure and develop a properly systematic understanding of the symbols. Jansen and Watter's ANN is a successful example of this, drawing on sensorimotor data to figure out word categories and syntactic formations and ultimately process language much more systematically. By capitalizing on the rich structure God built into our perceptual experience of the world, we can come much closer to having a computer maneuver both the flexibility of language that neural networks excel at and language's powerful logical structure that requires systematicity.