
Concepts and Chunks in Cognitive Systems

Pat Langley

PAT.LANGLEY@GTRI.GATECH.EDU

Information and Communications Laboratory, Georgia Tech Research Institute, Atlanta, GA 30318
Institute for the Study of Learning and Expertise, 2164 Staunton Court, Palo Alto, CA 94306

Abstract

Research in cognitive science, both empirical and theoretical, has implicitly distinguished between two types of mental structures: concepts and chunks. The first paradigm has studied people's ability to represent, recognize, and learn categories, with classification error and response time as typical measures. The second has examined humans' ability to store, use, and acquire recurring patterns, with memory capacity as the main dependent variable. The literatures on these two phenomena are nearly disjoint, but they also share many assumptions. In this essay, I review the similarities and differences between computational accounts of concepts and chunks, then propose some steps toward developing a unified theory of their roles in natural and artificial cognitive systems.

1. Introduction

The cognitive systems paradigm has a number of characteristics that distinguish it from other computational approaches to intelligence (Langley, 2012). Perhaps the most basic feature is its emphasis on the central role of *cognitive structures*. Of course, these alone do not suffice to explain mental abilities; they must also be organized in some manner, performance mechanisms must use them effectively, and learning processes must acquire and refine them. However, these all build on cognitive structures, so discussions naturally start with them.

Like other scientific disciplines, the cognitive systems movement aims for general theories with broad coverage of phenomena, in this case abilities found in human intelligence. Such coverage appears to require multiple types of mental structures and mechanisms that operate on them, as illustrated in many cognitive architectures (Langley, Laird, & Rogers, 2009). At the same time, we should not postulate different types of content unless necessary, as elegant theories are more desirable than complicated ones, other things being equal.

In this essay, I discuss the prospects for unifying two classic types of cognitive structures – *concepts* and *chunks* – that have typically been treated separately by cognitive psychologists and AI researchers. I review computational accounts of these two notions in turn, including common assumptions about representation, organization, performance, and learning. Next I summarize similarities and differences in these narratives, after which I propose an agenda for developing unified theories that would cover the phenomena associated with both types of structures. In closing, I outline three promising alternatives for tackling this unification challenge.



2. Concepts in Cognitive Systems

The study of concepts and categories has been a mainstay of research on cognitive psychology since the 1960s, as discussed by Rips, Smith, and Medin (2012), and it has played an equally important role in artificial intelligence, especially the subfield of machine learning (Langley, 1996). Here is one candidate definition:

- *A concept is a cognitive structure that denotes a set of entities or situations and that characterizes them in terms of their regularities.*

This statement is intentionally abstract, as needed if we want to cover the variety of accounts that have been proposed. However, all treatments share the assumptions that concepts involve mental structures, that these refer to some set or class, and that this set is not arbitrary but rather is characterized by law-like descriptions. The terms *category* and *concept* may be used synonymously, but the former often refers to the set being denoted and the latter to mental structures that encode it.

Humans associate words with many concepts in everyday life, most of them referring to classes of objects, whether inanimate (e.g., clouds, street signs) or animate (e.g., dogs, cats), but not all of them are linked to language. There can be considerable variation among the members of each category, yet people group them together because they nevertheless share underlying similarities. Specialists have their own concepts, from experts in automobile diagnosis who distinguish types of faults to professional wine tasters who discriminate among different vintages. Concepts help us encode, store, organize, and communicate our experience of the world.

If concepts are mental structures, then a cognitive system must *represent* them in some manner. The literature reflects many competing positions on this issue, but Medin and Coley (1998) have organized them into three broad classes.¹ These include:

- *Logical accounts*, in which a concept specifies the necessary and sufficient conditions for membership in a category. This view, perhaps the earliest one, has roots in philosophy. However, treating concepts as logical conjunctions appears too constraining for many human categories, and some variants allow for disjunctive descriptions.
- *Prototype accounts*, in which a concept provides a generalized description for a category that notes relevant features or relations but allows for differing degrees of membership, with some instances of a category fitting better than others. Probabilistic representations of concepts are a widely adopted instance of this idea.
- *Exemplar accounts*, in which a concept comprises a set of representative exemplars that, taken together, indicate membership in the category. As with prototypes, some instances of a category will fare better than other ones based on their similarity to the stored exemplars, with notions of similarity differing across alternative models.

Langley (1996) mentions a fourth option – *threshold concepts* – that include some neural network accounts. There are variations on each theoretical position, often introduced in response to observed phenomena in human behavior. There is not enough space to cover them all here, but the objective of this paper is not to review them fully, only to set the stage for a later comparison with chunks.

1. Medin and Coley refer to the first of these views about conceptual representation as *classical* and the second as *probabilistic*, but the labels *logical* and *prototype* also appear in the literature.

Some theories about concepts also make statements about the *organization* of these conceptual structures in long-term memory. The simplest, sometimes implicit, position is that memory simply contains a set of concepts, with no further constraints on their storage. Another account (e.g., Fisher, 1987) claims that concepts are organized in a taxonomic hierarchy with more general categories above their more specific descendants. A related but distinct alternative states that concepts serve as nodes in a lattice or directed acyclic graph, which means that they can have multiple parents. A third scheme organizes knowledge in an *inference network* (Langley, 1996), in which simpler concepts provide evidence for more complex ones, with neural networks and logic programs being examples. These different frameworks provide competing indexing schemes for conceptual structures.

As Anderson (1978) has argued, representational assumptions have little predictive or explanatory power without additional claims about how the structures are *used*. Thus, theories about conceptual knowledge typically come with postulates about performance mechanisms that operate over them. In particular, accounts that state concepts are logic-like in character typically posit an all-or-none matching process to determine whether they apply to given situations. In contrast, prototype theories suppose some form of partial matching that computes a similarity score or a probability for a concept description given a situation.² Finally, exemplar or case-based approaches apply the same basic idea to calculate similarities between stored cases and a situation. Typical measures of performance are classification error (or accuracy) and recognition (response) time.

There is also a substantial literature on concept *acquisition* in humans and machines. Work in both communities has concentrated on supervised training, in which each case has an associated class label, but some efforts have addressed unsupervised learning (e.g., Fisher et al., 1991). Research in machine learning has relied mainly on batch handling of training cases, but it is clear that people process experiences incrementally, so many models of human concept learning make this assumption. A variety of incremental mechanisms have appeared in the literature, including generalization from specific cases, specialization of overly general structures, and probabilistic updates of statistical descriptions. One characteristic of human acquisition that has received less attention is its piecemeal character (Langley, 2022), which involves learning one structure at a time.

3. Chunks in Cognitive Systems

The study of chunks has also received considerable attention in cognitive psychology (Miller, 1956; Fountain & Doyle, 2012), although somewhat less than concepts. The idea has also played a role in artificial intelligence, although seldom by that name, with Campbell and Berliner's (1983) work on chunks in chess being an important exception. Here is a tentative definition:

- *A chunk is a cognitive structure that denotes a collection of entities, possibly typed, that form a specified configuration or relational pattern.*

Again, this statement is intentionally abstract in order to cover the range of theories that have appeared in the literature. Nevertheless, despite their differences, these accounts share assumptions that chunks are cognitive structures, they refer to collections of entities, and they are characterized as patterns of relations among their constituents.

2. Note that theories of structural analogy (e.g., Gentner & Forbus, 2011) assume partial matching over logic-like structures, which supports the argument that performance mechanisms are crucial.

Everyday examples of chunks include written words, telephone numbers, familiar tunes, standard table settings, and the layout of rooms in a house. They also arise in more specialized context, and some of the strongest evidence for them comes from studies of expert memory. For instance, Chase and Simon (1973) found that chess masters store and use chunks for many board configurations. Similarly, it is obvious that chefs know the ingredients for many recipes, actors recall extended passages from many plays, and experienced taxi drivers remember the layout for their home cities. Chunks are as pervasive as concepts in human cognition.

Consider the above definition’s implications for mental representation. A chunk specifies a collection of entities that satisfy some relational pattern or configuration. This does not constrain the constituent entities, so they may be physical objects (e.g., silverware on a table), marks on a page (e.g., letters), or even complex activities (e.g., dance steps). Neither does it limit the nature of the relations, which may be spatial (as in constellations), temporal (as in music), or something more abstract. Moreover, these relations may be purely qualitative (e.g., above, behind) or they may include quantitative details (e.g., distance, angle). Finally, the literature usually treats chunks as logic-like structures, but nothing prevents them from including probabilities or other annotations.

The definition also allows chunk constituents to have types, which means the latter may themselves be other chunks. This has import for memory organization, as it implies that a chunk can specify a hierarchical structure. However, this does not involve the taxonomic hierarchies that organize some concepts, but rather a partonomic or compositional hierarchy. For example, a phrase can be broken down into a sequence of words, each of which comprises a sequence of letters. Not all treatments of chunks make this assumption, but they do not rule it out and there is general agreement that high-level chunks may be specified in terms of lower-level ones.

Theories of chunk use in humans emphasize two primary capabilities. The first involves the *recognition* of chunks from perceptions, such as a friend’s telephone number, a familiar passage of music, or a famous constellation in the night sky. Accounts in the literature typically assume this is a bottom-up, all-or-none match process, with relations among elements playing a key role. The second involves *recall* or reconstruction of chunks, typically after some delay. Examples here include remembering a sequence of digits or words after hearing them or recreating positions on a chess board after a brief exposure. In this case, most accounts assume some form of top-down generative process. Typical measures here are memory capacity and reconstruction time.

There have been some computational accounts of chunk acquisition in humans, but, again, fewer than for concept learning. Computational models like EPAM (Richman, 1991) and CHREST (Gobet & Lane, 2012) assume, based on empirical evidence, that this process is incremental and largely unsupervised, with mastering the ability to recognize chunks preceding the ability to reconstruct them. Educational studies also suggest that it is typically cumulative, with high-level chunks being acquired only after those for their constituents have been learned. Research in machine learning on tasks like grammar induction (e.g., Wolff, 1982; Langley & Stromsten, 2000) has often used nonincremental methods to introduce chunks, but these appear to differ from those used by humans.³

3. The ‘chunks’ acquired by the Soar architecture (Laird, 2012), which primarily encode procedural knowledge, bear little resemblance to Miller’s original notion, which emphasized perception and memory.

4. Analysis of Similarities and Differences

Now that we have reviewed key ideas from the literature on concepts and chunks, we can compare and contrast them. As seen above, their study has emphasized different aspects of cognition, but this does rule out the possibility that one type of mental structure, or even a single set of processes, underlies both sets of cognitive phenomena. Because this is our main hypothesis, we should start by examining their common assumptions.

We can identify five primary postulates that are shared by the majority of psychological research on concepts and chunks. Many of these tenets also hold for their treatment in artificial intelligence. They include:

- *Discrete cognitive structures.* One common assumption is that concepts and chunks are both encoded as discrete structures which are stored in long-term memory. A given concept or chunk may be linked to others, but they remain separate mental entities with their own content.
- *Organization in memory.* Another shared postulate is that both concepts and chunks are organized and indexed in long-term memory. The forms of their organization differ, but discrete structures of a given type are nevertheless linked in some fashion.
- *Access through recognition.* Theories of concepts and chunks both posit they can be accessed from long-term memory through a recognition process. This involves matching their descriptions against perceptions, but it may also take advantage of their organization.
- *Response time as a performance metric.* Empirical studies of concepts and chunks in humans often use response time as a measure of performance. This involves recognition time for concepts and reconstruction time for chunks, but both equate rapid processing with mastery.
- *Incremental and piecemeal learning.* Studies of human learning reveal that it occurs incrementally, with one or a few instances being processed at a time, and in a piecemeal manner, with one or a few cognitive structures being acquired or revised at a time.

These commonalities offer some encouragement that a unified account of concepts and chunks is possible, and efforts toward this goal can use them as constraints on candidate theories, as already done in many separate accounts.

But as we have seen, despite their underlying similarities, the research community has generally treated concepts and chunks as though they are different types of mental entities. Some important distinctions between them include:

- *Categories vs. composites.* Theoretical treatments of concepts stress their ability to represent sets of instances or categories, whereas accounts of chunks emphasize their compositional character. Chunks may cover different instances, but this is downplayed. Concepts may involve relations, but this is not required, whereas relational patterns are central to chunks.
- *Taxonomic vs. partonomic hierarchies.* Not all theories of concepts discuss organization, but a common idea is that they are positioned in a taxonomic hierarchy, with higher nodes being more general than their descendants. In contrast, theories of chunks typically posit they are organized in a partonomic hierarchy, with parents being composed of their children.
- *Flexible vs. strict matches.* Most accounts of conceptual processing assume flexible matching, with classic examples being prototype and exemplar theories. Explanations of chunk retrieval often treat their defining patterns as abstract yet rigid, with relations standing or falling together.

Table 1. Different emphases in the treatments of concepts and chunks.

Dimension	Concepts	Chunks
Representation	Categories	Composites
Organization	Taxonomic	Partonomic
Matching	Flexible	Strict
Performance	Recognition	Reconstruction
Acquisition	Generalization	Composition

- *Recognition vs. reconstruction.* In most cases, studies of conceptual knowledge emphasize their ability to *recognize* instances of categories, whereas treatments of chunks usually emphasize their ability to *reconstruct* patterns from memory.⁴ Some accounts of concepts discuss their generative capacity, and chunks are certainly recognized, but these tend to be downplayed.
- *Generalization vs. composition.* Treatments of concept acquisition emphasize generalization over training cases. Learning occurs incrementally with each instance, but multiple examples are often needed to establish category bounds. Theories of chunk acquisition focus on composition of elements into larger structures. Mastery may require multiple trials because elements may be added incrementally. Moreover, chunking supports cumulative learning, in which later structures build on earlier ones, an idea far less common with concepts.

Table 1 summarizes these differences, but it is important to note they are primarily distinctions in emphasis. Theories of conceptual knowledge seldom forbid the features associated with chunks and theories of chunks implicitly allow those associated with concepts. This leaves the door open for an account that combines their features, but we must still walk through this intellectual passage.

5. A Proposed Research Agenda

Although a computational explanation of human-like intelligence that treats concepts and chunks separately may be possible, it would be considerably less satisfying than a unified theory. There have been a few forays in the latter direction, but such attempts have received limited attention and the topic merits further effort. In this section I consider some steps that the cognitive systems community might take toward this end.

5.1 Representation

As noted, theories of concepts emphasize the representation of generic categories or sets of items, whereas those for chunks instead focus on their compositional character. A unified account would incorporate both ideas by positing a single type of cognitive structure that describes a relational pattern of elements while also characterizing a class of instances that satisfy it. Such structures

4. The distinction between recognition and recall in cognitive psychology maps roughly onto the dichotomy between *discriminative* and *generative* models in machine learning.

would have aspects of both concepts and chunks, although features of one of the other would still be dominant in degenerate cases.

For instance, a written word would be specified as a sequential pattern of letters, but that word might have alternative spellings which provide generality. Similarly, a structure for a chess board would involve a spatial pattern of pieces that threaten or defend each other, but different types of pieces might occupy the same role, denoting a set of possible configurations. One could describe a passage of music as a sequence of notes, but this could allow variations in timing and emphasis in how it is played, again covering a range of instances.

Of course, a full theory would need to specify the details of this unified representation. In one version, some structures would be more concept-like, with little compositional character, whereas others would be more chunk-like, with little room for variation. In another version, every structure would reflect both aspects, exhibiting both the generality of concepts and the configurational nature of chunks. Whether these specifications are best encoded as logical descriptions, prototypes, or exemplars remains an open question.

5.2 Organization

We have seen researchers often assume that concepts are organized in a taxonomic hierarchy, while chunks instead participate in partonomic arrangements. A unified theory would incorporate both forms of mental organization, with structures that refer to constituents and thus impose a part-of hierarchy, but nevertheless reside within a taxonomy in which higher-level nodes are generalizations of their descendants. Human cognition supports both forms of organization and our computational theory of intelligence should do so as well.

For example, the structure for mammals would include features like hair covering and placentas, but it would also include constituents like a head, torso, and limbs in a spatial configuration. Different categories of mammals, such as quadrupeds and bipeds, would be specializations that reside lower in the taxonomy. However, the latter would also contain nodes for categories of constituents, such as hooved and clawed limbs, that appear in their concept-like descriptions. Similarly, word classes would be specified as sequences of letters, with some variation, which in turn would be described by their arrangement, with greater variation due to handwriting style or printed fonts.

A unified theory would specify how these two forms of organization interleave and complement one another. For instance, each node in a taxonomic hierarchy could specify constituents that make up that concept, but these components would also reside in the same taxonomy or a kindred one. The account should also state whether the relations between chunk constituents are themselves decomposable or whether they allow variation across instances. Formalisms like context-free grammars and logic programs have some of the required features, but more appear necessary.

5.3 Performance

As noted above, many treatments of concept use favor flexible forms of recognition, while accounts of chunk use emphasize recall or reconstruction. A unified theory would incorporate both capabilities, relying on a form of partial matching when comparing stored structures to perceptions and drawing on their relational patterns for reconstruction from long-term memory. Most theories of

concepts and chunks do not rule this out, but their focus on different phenomena means they have each told only part of the overall story.

For instance, a bottom-up recognition process could group strings of letters into words and then words into phrases to construct parse trees for sentences, producing the same hierarchical structures for different fonts and spellings. In addition, a recall process could fill in missing letters or generate words for incomplete sentences. The same mechanisms could recognize familiar configurations in chess boards, even if they involve novel pieces, and reconstruct their layout on the board from a chunk identifier and its constituent entities.

In a unified account, structures would be recognized by finding their degree of match to perceptions, possibly by sorting them through a taxonomic hierarchy that indexes them. This process might occur in a bottom-up manner, with categorization of constituents like letters leading in turn to recognition of composites like words, but top-down forms of retrieval might also be possible. Reconstruction would involve the decomposition of structures into constituents, possibly through multiple levels, using stored relations to organize the generative process.

5.4 Acquisition

We have seen that theories of concept learning emphasize generalization over experience, whereas accounts of chunking focus on creation of new relational structures composed of elements. A unified framework would incorporate both forms of acquisition, letting it create generalized descriptions that cover distinct training cases but that also specify relations among their constituents. This need not rely on a single learning mechanism, but if it involves multiple processes, they must work in tandem to support both facets of acquisition.

For example, this cooperative arrangement would support learning of letter ‘concepts’, which can vary substantially in appearance, but also construction of word ‘chunks’ with letters as components. However, different words might occur in similar contexts, leading to ‘concepts’ like nouns and adjectives, which can then be combined into ‘chunks’ for adjectival phrases and noun phrases. Such conjoined processes would support cumulative learning in which lower levels of the partonomic hierarchy (each embedded in a taxonomy) would provide elements to compose higher levels.

A unified theory of concepts and chunks would explain how to acquire these two-faceted structures. A crucial question is whether two learning mechanisms are necessary, as in research on the induction of context-free grammars (e.g., Wolff, 1982; Langley & Stromsten, 2000), or whether a single process can suffice. The extended framework must acquire both concept-like and chunk-like features of new structures, but this might result from interactions between one learning process and multiple performance mechanisms.

6. Three Candidate Theories

We can explore the prospects for a unified theory by examining existing frameworks that appear to hold potential for handling both concepts and chunks. In this section, I review three candidate’s assumptions about representation, organization, performance, and learning, then discuss how researchers might extend their coverage to both sets of phenomena.

6.1 Discrimination Networks

One of the earliest computational models of human expertise was EPAM (Feigenbaum, 1961; Feigenbaum & Simon, 1984), which organized knowledge in a discrimination network. This takes the form of a tree, with each node indicating a test or attribute and each downward branch specifying an outcome or value, much like a decision tree (Quinlan, 1986). Each terminal node stores not a class label but an ‘image’ or description (e.g., a conjunction of attribute values). Performance involves sorting a stimulus downward through the network until reaching a terminal node and using its image to recall unspecified values. Learning is interleaved with performance, with a *discrimination* process adding a new branch when the stimulus matches no branch’s value and a *familiarization* mechanism adding details to an image on reaching a terminal node.

The earliest versions of EPAM focused on rote memorization, as required by tasks like paired-associates learning. However, later ones (e.g., Richman, 1991) augmented the framework to let it store images at internal nodes, making their discrimination networks similar to taxonomies of concepts. CHREST (Gobet & Langley, 2012; Lane, Gobet, & Smith, 2008) further extended these images into ‘templates’ whose roles can be filled in different ways, making them even more like generalized concepts. Both Richman and Lane et al. claimed that their discrimination networks supported chunk storage, use, and acquisition, but they seem to have equated chunks with images, and did not provide a way to acquire high-level chunks in terms of simpler ones.

We should consider what must be added to support the compositional character of chunks and their cumulative acquisition. EPAM’s tests were not limited to primitive features; they could refer to nodes elsewhere in the network. This meant that it could define complex chunks (e.g., words as sequences of letters) in terms of simpler ones (e.g., letters as configurations of lines). However, in such cases, EPAM and its successors appear to have required stimuli that were already organized into partonomic structures. What they lacked was some means to acquire these structures from training data and some way to infer them during performance. Techniques from grammar induction (e.g., Wolff, 1982) that introduce nonterminal symbols offer one response to the first issue and classic methods for parsing offer an approach to the second. Integrating these with the framework of discrimination networks could offer a unified account of concepts and chunks.

6.2 Neural Networks

A second promising paradigm, multilayer neural networks, organizes expertise in a very different manner. Nodes denote features or attributes and weighted links connect them to others, typically in a directed acyclic graph. Input nodes at the lowest level feed into intermediate ones, often in many layers, leading ultimately to output nodes at the highest level. Performance involves passing values of perceived features to input nodes, using the weighted links to compute values for intermediates, and finally calculating values for output nodes. For classification tasks, the system typically selects the highest-scoring node to assign a class. Learning usually relies on some variant of backpropagation (Rumelhart et al., 1986), which alters weights on links to reduce errors in output values.

Because neural networks are widely used for classification tasks, they have clear relevance to concepts, and Kruschke’s (1992) ALCOVE was an early model of human categorization cast in these terms. However, we can also view their internal nodes as encoding chunk-like structures, although this is not readily apparent because they take on continuous values, whereas classic chunks

are discrete. Yet on visual processing tasks like letter and face recognition, examination of the weights on hidden units show they correspond to constituent features for edges, corners, and other facets of images. This suggests that the widespread claim neural networks exhibit *representation learning* by creating latent variables is equivalent to stating they acquire chunks.

Based on this observation, we might conclude that neural networks already provide a unified account of concepts and chunks, even though this idea has not received attention in the literature. The fact that they must be initialized with a complete network rather than adding nodes (chunks) as needed is not an obstacle, as starting weights on links to intermediate nodes can be negligible. The problem lies not with representation or performance, but with learning, as humans acquire chunks and concepts far more rapidly than the gradient descent methods that neural networks use to update their parameters. To support human-like behavior, the framework will need alternative mechanisms that learn in a more rapid and piecemeal manner (Langley, 2022).

6.3 Probabilistic Concept Hierarchies

A third theoretical framework, first introduced in Cobweb (Fisher, 1987; Iba & Langley, 2011), focused explicitly on concept formation. This assumes that concepts are organized as nodes in a taxonomic hierarchy, with terminal nodes denoting training cases and nonterminals describing probabilistic summaries of their descendants. A performance mechanism sorts new cases down the path that maximizes an information-theoretic metric, predicting missing features as needed. Learning is incremental and unsupervised, with each node's distribution being updated as instances pass through it and new concepts being created when reaching a terminal node or lacking a good-scoring child. The framework combines ideas from discrimination network, probabilistic, and exemplar accounts, and it exhibits behavior consistent with some categorization phenomena.

The variants of Cobweb developed to date have focused on conceptual knowledge, but some extensions could let the framework support chunk representation, organization, use, and acquisition as well. To this end, they might assume that instances comprise a set of elements (e.g., words or objects), each with a description but also linked by relations (e.g., order or direction), and that concepts specify a set of constituents and their surrounding contexts, along with the probability distributions for each. They might also include a parsing mechanism that, when processing an instance, creates candidate chunks that it sorts through the taxonomy to select the best, which it then uses to elaborate the instance. Every nonprimitive concept would have constituents and thus also be a chunk. Parsing would use this knowledge to recode perceived instances by adding partonomic structure. Learning would remain largely the same except for operating over the extended representation.

The commitment that all but primitive training cases contain multiple elements would ensure the constituents that provide material for parsing and chunk creation. The inclusion of context would allow creation of concepts based on this information rather than ones based only on similar content. This would let Cobweb acquire grammatical knowledge by identifying words (black, cat) as letter sequences, word classes (adjective, noun) in terms of shared contexts, phrasal structures as sequences of word classes, and phrasal classes that arise in shared contexts at a higher level. Presumably, the same mechanisms could explain chunk processing in domains like chess. Chunks in Cobweb would include probabilities, but highly regular ones would behave like logical structures, as in classic accounts. The resulting theory would unify concepts and chunks.

7. Summary Remarks

In this paper, I noted the cognitive system paradigm's emphasis on cognitive structures and examined two varieties – concepts and chunks – that have received attention in cognitive psychology. I reviewed common assumptions about the representation, organization, use, and acquisition of concepts and chunks, then examined similarities and differences in their treatments. I concluded that these two types of entities may not be truly distinct, but rather simply facets of the same mental encodings viewed from different perspectives.

In response, I proposed a research agenda for developing unified theories of concepts and chunks. This would explore generalized representations that subsume both ideas, common organizations for them in long-term memory, unified performance mechanisms that operate over them, and learning processes that acquire and refine them. I also discussed three existing frameworks that seem ripe for extension. We should evaluate each in terms of its ability to support abilities associated with both concepts and chunks, subject to constraints on human cognition. Progress in this direction would take us an important step closer to a unified theory of the mind.

Acknowledgements

Work on this essay was supported by Grant FA9550-23-1-0580 from the US Air Force Office of Scientific Research and by Grant N00014-23-1-2525 from the Office of Naval Research, which are not responsible for its contents. I thank Chris MacLellan, Doug Fisher, and many others for discussions that contributed to the ideas presented here.

References

- Anderson, J. R. (1978). Arguments concerning representations for mental imagery. *Psychological Review*, 85, 249–277.
- Campbell, M., & Berliner, H. J. (1983). A chess program that chunks. *Proceedings of the Third National Conference on Artificial Intelligence* (pp. 49–53). Washington, DC: AAAI Press.
- Chase, W. G., & Simon, H. A. (1973). Perception in chess. *Cognitive Psychology*, 4, 55–81.
- Feigenbaum, E. A. (1961). The simulation of verbal learning behavior. *Proceedings of the Western Joint Computer Conference* (pp. 121–132). Los Angeles, CA.
- Feigenbaum, E. A., & Simon, H. A. (1984). EPAM-like models of recognition and learning. *Cognitive Science*, 8, 305–336.
- Fisher, D. H. (1987). Knowledge acquisition via incremental conceptual clustering. *Machine Learning*, 2, 139–172.
- Fisher, D. H., Pazzani, M. J., & Langley, P. (Eds.) (1991). *Concept formation: Knowledge and experience in unsupervised learning*. San Francisco, CA: Morgan Kaufmann.
- Fountain, S. B., & Doyle, K. E. (2012). Learning by chunking. In N. M. Seel (Ed.), *Encyclopedia of the sciences of learning*. Springer, Boston, MA.
- Gentner, D., & Forbus, K. D. (2011). Computational models of analogy. *Wiley Interdisciplinary Reviews: Cognitive Science*, 2, 266–276.
- Gobet, F., & Lane, P. C. R. (2012). Chunking mechanisms and learning. In N. M. Seel (Ed.), *Encyclopedia of the sciences of learning*. New York, NY: Springer.

- Iba, W. F., & Langley, P. (2011). Cobweb models of categorization and probabilistic concept formation. In E. M. Pothos & A. J. Willis (Eds.), *Formal approaches in categorization*. Cambridge: Cambridge University Press.
- Kruschke, J. K. (1992). ALCOVE: An exemplar-based connectionist model of category learning. *Psychological Review*, 99, 22–44.
- Laird, J. E. (2012). *The Soar cognitive architecture*. Cambridge, MA: MIT Press.
- Lane, P. C. R., Gobet, F., & Smith, R. L. (2008). Attention mechanisms in the CHREST cognitive architecture. *Proceedings of the Fifth International Workshop on Attention in Cognitive Systems* (pp. 207–220). Graz, Austria: Joanneum Research.
- Langley, P. (2012). The cognitive systems paradigm. *Advances in Cognitive Systems*, 1, 3–13.
- Langley, P. (2022). The computational gauntlet of human-like learning. *Proceedings of the Thirty-Sixth AAAI Conference on Artificial Intelligence* (pp. 12268–12273). Vancouver: AAAI Press.
- Langley, P., Laird, J. E., & Rogers, S. (2009). Cognitive architectures: Research issues and challenges. *Cognitive Systems Research*, 10, 141–160.
- Langley, P., & Stromsten, S. (2000). Learning context-free grammars with a simplicity bias. *Proceedings of the Eleventh European Conference on Machine Learning* (pp. 220–228). Barcelona, Spain: Springer-Verlag.
- Rips, L. J., Smith, E. E., & Medin, D. L. (2012). Concepts and categories: Memory, meaning, and metaphysics. In K. J. Holyoak & R. G. Morrison (Eds.), *The Oxford handbook of thinking and reasoning*, 177–209. Oxford, UK: Oxford University Press.
- Medin, D. L., & Coley, J. D. (1998). Concepts and categorization. In J. Hochberg (Ed.), *Perception and cognition at century's end*, 403–439. New York, NY: Academic Press.
- Miller, G. A. (1956). The magical number seven, plus or minus two: Some limits on our capacity for processing information. *Psychological Review*, 63, 81–97.
- Quinlan, J. R. (1986). Induction of decision trees. *Machine Learning*, 1, 81–106.
- Richman, H. B. (1991). Discrimination net models of concept formation. *Concept formation: Knowledge and experience in unsupervised learning*. San Francisco, CA: Morgan Kaufmann.
- Rumelhart, D. E., Hinton, G. E., & Williams, R. J. (1986). Learning internal representations by back-propagating errors. In D. E. Rumelhart & J. L. McClelland (Eds.), *Parallel distributed processing* (Vol. 1), 318–362. Cambridge, MA: MIT Press.
- Wolff, J. G. (1982). Language acquisition, data compression and generalization. *Language & Communication*, 2, 57–89.