Schematic processing as a framework for learning and creativity in CBR and CC

Kat Agres and Geraint A. Wiggins

Queen Mary University of London, London E1 4FZ, United Kingdom,
{`kathleen.agres`,`geraint.wiggins`}@`qmul.ac.uk`

Abstract. There is a clear connection to be made between psychological findings regarding learning and memory and the areas of case-based reasoning (CBR) and computational creativity (CC). This paper aims to encourage researchers in these areas to consider psychological perspectives while developing the technical and theoretical aspects of their computational systems. To this end, an overview of knowledge structures and schematic processing is provided, offering findings from music cognition to demonstrate the utility of this approach. Examples of musical schemata are offered as cases which may be used in CBR systems for combinatorial creativity and the generation of new creative output.

Keywords: cognitive psychology, schematic processing, computational creativity, case-based reasoning

1 Introduction

Creativity relies heavily upon domain-relevant experience and knowledge: an expert chess player's creative problem-solving, for example, is based on his robust knowledge and flexible thinking within his domain. Given the prime importance of past learning and experience for future creative behavior, there is an obvious marriage between the areas of case-based reasoning (CBR) and computational creativity (CC). While this connection has been explored in various computational settings, few approaches import findings and perspectives from cognitive psychology (although, see [10]), a field which may offer rich insight into this endeavour. Specifically, the mechanisms underlying learning and memory, and the way in which information is represented in the mind, should be considered, as these can elucidate creative behavior and inspire new ways of approaching machine creativity. In other words, artificial systems simulating human learning and memory can form the foundation for CBR approaches to CC.

This paper takes the stance that considering psychological mechanisms is essential not only for understanding human creativity, but for a theoretical understanding of creativity that can inform the implementation of creative processes in artificial systems. That is, researchers may be able to bolster CC by understanding how humans are creative. We focus on schematic processing mechanisms, such as the encoding and updating of memory representations, and the domain of music is considered as an example of how the abstraction of instances or cases yield schemata (e.g., generalized cases) which may be applied to CC.

2 Knowledge structures in human cognition

Cognitive psychology has thoroughly investigated learning and memory. Researchers once believed memory to be vast and detailed [28], but recent findings highlight its incompleteness and malleability. For example, vision research suggests that viewers primarily encode the general schematic attributes of a visual scene upon brief initial viewing [20, 26], supplying a semantic understanding of the scene [19, 20] but lacking detail. Similary, psychology and cognitive science have recently emphasized the importance of association and analogical processing [1,11]. Although veridical representations are sometimes encoded, more often we form general or associative semantic representations (*schemata*) of new input based on prior experience. This *schematic processing* is based on abstracted mental representations that structure or organize some aspect of past experience, and schematic memory structures influence the processing of new information.

Investigations of schematic processing have contributed to our theories of learning and memory for nearly a century [2, 24]. In *Remembering*, Bartlett notes that when individuals are asked to recall an odd or supernatural story after a time delay, their recollections alter the story to better conform to their existing schematic knowledge [2]. In other words, our knowledge shapes our perception and interpretation of the world. Piaget, who considered schemata to be the building blocks of knowledge, discussed how new information is incorporated into existing schemata in the processes of *assimilation* [24]. When the new information is too dissimilar to be integrated, *accommodation* occurs, in which the schematic structure itself must change to accommodate the new information.

The notion of schemata has been echoed in the fields of computer science and artificial intelligence for decades, for example, in Minsky's *frames* [17], and Schank's *script-based systems* [27]. Recent computational models learn and generalize the statistics of a training corpus (building what is essentially a statistical version of a schematic framework) in order to evaluate or categorize new instances [13, 23]. This is akin to the process of assimilating new information into schematic representations, where the schemata in this instance are encoded in the network of probabilities underlying common structures or patterns. These statistical models have been used to generate new, creative output [22, 25]. CBR and CC approaches have successfully used techniques such as inductive analogical processes [21] and template-based methods (e.g., Gervas' ASPERA system [8]) for creative generation, but the connection to schema theory is often only implicit. Arguably, psychological findings should be explicitly applied here, because knowledge of how mental representations are formed and change over time (and are re-represented) can inform how AI systems may represent the information and knowledge required to achieve creative behaviors.

3 Music as an example domain

To show how psychology can inform how systems learn, represent, and combine information in new ways, we consider the domain of music. In the auditory modality, Bregman, Dowling, Cuddy, and others have explored the contribution of schema-based mechanisms to the abstraction of tonal relationships during music perception [4, 5]. Experience listening to common musical patterns or forms creates our mental framework for processing music [9, 16]. The underlying schemata are essentially collections of rules that guide listeners' perception of music (and thence the information encoded) by directing attention and continually creating expectations about the forthcoming music [12, 15, 18, 23]. Although musicians may have more elaborated schemata than non-musicians, everyone exposed to music has implicitly learned musical schemata. Conversely, every schema is modified by perceptual experience, as new information is abstracted and integrated into long-term schematic memory [29].

For concrete examples of musical schemata, we may consider Gjerdingen's examples of musical schemata: the "gap-fill" schema and the "changing-note" schema [9]. The former matches a melodic leap followed by an ascending or descending sequence of tones that fills the gap created by the interval leap. The latter matches two pairs of notes, in which the first pair leads away from the tonic pitch, and the second leads back. Even musically untrained listeners are capable of distilling these schemata from examples containing both types [9]. He further argues that musical schemata comprise a specific set of features that create a *style structure* [18]. Similarly, Snyder [29] describes musical schemata as networks of long-term memory associations that are amalgamations of the statistical properties of music: semantic frameworks constructed from "the commonalities shared by different experiences" [29]. Over time, episodic memories gradually form a generalized schematic representation in which specific details of each instance are lost, but generalizability of the schemata is gained.

In sum, musical schemata are mental frameworks of musical knowledge that are abstracted from experience and guide musical expectation. One insight from this work for CBR is to not simply match cases, but to *generalise* cases into schemata. If a CBR system has internalized schemata based on a corpus of musical cases (e.g., melodies), it is equipped to process new examples with more sophistication: by extracting schematic representations of these melodies, the representations may be more easily compared, and the generation of new music is made more feasible. Consider a system that generates novel, high-quality harmonization. First, it is provided with a case base of well-harmonized melodies from which it extracts schemata and derives characteristics of good harmonization. Then, given a new melody (case), it can generate harmony by matching within the space of schemata, to extrapolate a novel but appropriate harmony.

4 Knowledge structures as the foundation for creativity

Learning mechanisms and knowledge representations (such as schemata) are essential to how humans structure and combine information. They are also of central importance to CC, and the principle of combining existing knowledge into novel ideas has been a cornerstone of creativity research for decades [3, 6, 14]. Koestler describes creativity as *bisociation*—"interlocking of two previously unrelated skills, or matrices of thought" [14]. Inspired by Koestler, Fauconnier and Turner [6] offer a cognitive theory of conceptual blending, in which elements and relationships from different sources are combined to produce new meaning.

Several authors also refer to conceptual spaces which may be combined, manipulated, and traversed [3, 7, 30]. In all of these approaches, schemata could be used as general cases (or matrices or regions of conceptual spaces) that may be combined to form new, creative ideas. Further, schemata may be viewed as methods for caching or even hashing the case base, thus improving retrieval efficiency.

Knowledge of psychological processes can inform how learning and memory may be instantiated in artificial systems, which in turn influences how concepts may be blended and combined. One may consider schemata to be the building blocks for exploratory and combinatorial creativity. If a CBR system maps melodic onto schematic representations, the system may then be used to classify or even generate new examples through extrapolation (or interpolation) of existing cases. This approach is especially useful for CC, because a means of reflection or self-evaluation should be built into the system, and CBR can satisfy this need. Further, the way in which humans learn and encode information can suggest particular schemata that may contribute to CC in AI systems, but also (and just as importantly), elucidate the *processes* underlying the combination of knowledge structures [30]. For example, one could use a schema-based system to judge whether new melodies will sound novel to listeners by examining whether different melodies abstract to the same schemata, and this could be very useful for applications such as automatic composition.

5 Conclusion

We argued for the consideration and inclusion of psychological findings in CBR as a means of approaching CC. Using examples of mental knowledge structures and schematic processing mechanisms in the musical domain, we discussed how existing schemata may be considered as cases for the combination of ideas and generation of new creative output. Understanding how humans learn and form memory representations may inform machine learning and CBR techniques, and ultimately, the expression of creativity in artificial systems.

Acknowledgements

The Lrn2Cre8 project acknowledges financial support of the Future and Emerging Technologies (FET) programme within the Seventh Framework Programme for Research of the European Commission, under FET grant number 610859.

References

1. Bar, M.: The proactive brain: using analogies and associations to generate predictions. Trends in cognitive sciences 11(7), 280–289 (2007)
2. Bartlett, F.C.: Remembering: A study in experimental and social psychology, vol. 14. Cambridge University Press (1995)
3. Boden, M.A.: The creative mind: Myths and mechanisms. Psychology Press (2004)
4. Bregman, A.S.: Auditory scene analysis: The perceptual organization of sound. MIT press (1994)

5. Dowling, W.J.: Scale and contour: Two components of a theory of memory for melodies. Psychological review 85(4), 341 (1978)
6. Fauconnier, G., Turner, M.: The way we think: Conceptual blending and the mind's hidden complexities. Basic Books (2008)
7. Gärdenfors, P.: Conceptual spaces: The geometry of thought. MIT press (2004)
8. Gervás, P.: An expert system for the composition of formal spanish poetry. Knowledge-Based Systems 14(3), 181–188 (2001)
9. Gjerdingen, R.O.: A classic turn of phrase: Music and the psychology of convention. Univ of Pennsylvania Press (1988)
10. Heath, D., Dennis, A., Ventura, D.: Imagining imagination: A computational framework using associative memory models and vector space models. In: Proc. of the 6th International Conference on Computational Creativity. p. 244 (2015)
11. Hofstadter, D.R.: Analogy as the core of cognition. Keith J. Holyoak, and Boicho N. Kokinov. Cambridge MA: The MIT Press/Bradford Book (2001)
12. Huron, D.B.: Sweet anticipation: Music and the psychology of expectation. MIT press (2006)
13. Kingma, D.P., Welling, M.: Auto-encoding variational bayes. In: Proceedings of the International Conference on Learning Representations (ICLR), (2014)
14. Koestler, A.: The act of creation. Macmillan (1964)
15. Krumhansl, C.L.: Cognitive foundations of musical pitch, vol. 17. Oxford University Press New York (1990)
16. Lerdahl, F.: Tonal pitch space. Oxford University Press (2001)
17. Minsky, M.: A framework for representing knowledge. Tech. rep., Cambridge, MA, USA (1974)
18. Narmour, E.: The analysis and cognition of melodic complexity: The implication-realization model. University of Chicago Press (1992)
19. Oliva, A.: Gist of the scene. Neurobiology of attention 696(64), 251–258 (2005)
20. Oliva, A., Torralba, A.: Building the gist of a scene: The role of global image features in recognition. Progress in brain research 155, 23–36 (2006)
21. Ontañón, S., Plaza, E.: Amalgams: A formal approach for combining multiple case solutions. In: Case-Based Reasoning. Research and Development, pp. 257–271. Springer (2010)
22. Pearce, M.T., Wiggins, G.A.: Evaluating cognitive models of musical composition. In: Cardoso, A., Wiggins, G.A. (eds.) Proc. of the 4th International Joint Workshop on Computational Creativity. pp. 73–80. London (2007)
23. Pearce, M.T., Wiggins, G.A.: Auditory expectation: the information dynamics of music perception and cognition. Topics in cognitive science 4(4), 625–652 (2012)
24. Piaget, J., Cook, M., Norton, W.: The origins of intelligence in children, vol. 8. International Universities Press New York (1952)
25. Ponsford, D., Wiggins, G.A., Mellish, C.: Statistical learning of harmonic movement. Journal of New Music Research 28(2), 150–177 (1999), http://www.soi.city.ac.uk/ geraint/papers/JNMR97.pdf
26. Rensink, R.A., O'Regan, J.K., Clark, J.J.: To see or not to see: The need for attention to perceive changes in scenes. Psychological science 8(5), 368–373 (1997)
27. Schank, R.C., Abelson, R.P.: Scripts, plans, goals, and understanding: An inquiry into human knowledge structures. Psychology Press (2013)
28. Shepard, R.N.: Recognition memory for words, sentences, and pictures. Journal of verbal Learning and verbal Behavior 6(1), 156–163 (1967)
29. Snyder, B.: Music and memory: An introduction. MIT press (2000)
30. Wiggins, G.A.: A preliminary framework for description, analysis and comparison of creative systems. Knowledge-Based Systems 19(7), 449–458 (2006)