On Bringing Case-Based Reasoning Methodology to Deep Learning

David Leake and David Crandall

Luddy School of Informatics, Computing, and Engineering, Indiana University
Bloomington IN 47408, USA
`leake@indiana.edu, djcran@indiana.edu`

Abstract. The case-based reasoning community is successfully pursuing multiple approaches for applying deep learning methods to advance case-based reasoning. This "Challenges and Promises" paper argues for a complementary endeavor: pursuing ways that the case-based reasoning methodology can advance deep learning. Starting from challenges in deep learning and proposed neural-symbolic integrations based on specific technologies, it proposes studying how CBR ideas can inform choices of components for a new reasoning pipeline.

Key words: Automated Machine Learning, Case-Based Reasoning Methodology, Challenge Problems, Deep Learning, Integrations, Pipelines

1 Introduction

Recent years have seen great accomplishments in deep learning. These have led to enthusiasm in the case-based reasoning (CBR) [9] community for studying how to apply deep learning methods in service of case-based reasoning. For example, the Call for Papers for the *2019 Workshop on Case-based Reasoning and Deep Learning* states that the "successes of DL call for novel methods and techniques that exploit DL for the benefit of CBR systems."[1] Research presented at that workshop and other venues supports the promise of this approach for advancing case-based reasoning. This "Challenges and Promises" paper proposes that the CBR community consider the reverse perspective: How application of case-based reasoning can shape the design of deep learning systems and help to address challenges for deep learning and machine learning as a whole.

As background to the challenges, the paper begins by highlighting two views on questions to address to advance deep learning and AI as a whole, presented in invited talks by Yann LeCun and Henry Kautz at the AAAI 2020 Conference on Artificial Intelligence in New York, NY. These focus, respectively, on challenges for deep learning and architectures for integrating neural and symbolic methods.

The proposed integrations of neural and symbolic methods view each as a different technology, with particular strengths for particular types of tasks. We propose that the CBR community develop integrations shaped by a different

---

[1] https://iccbr2019.com/workshops/case-based-reasoning-and-deep-learning/

perspective, that of Ian Watson's treatment of CBR as a *methodology* [16]. In that view, case-based reasoning is seen as a general high-level process that defines a set of tasks, but for which the needed functionality can be implemented using various technologies, both neurally inspired and symbolic.

This perspective suggests an opportunity for the CBR methodology to shape the high-level design of component-based deep learning systems, with collections of subparts corresponding to components of the CBR process—retrieve, reuse, revise, and retain [1]—and encoding the CBR knowledge containers—vocabulary, case knowledge, similarity knowledge, and adaptation knowledge [12]. It also raises questions of how such components can be implemented and integrated. Especially interesting is the addition of forms of *case adaptation* in deep learning frameworks, to enable the transformation of solutions for novel contexts. CBR can also play an important role in automated machine learning (AutoML), by helping to exploit experiences with AutoML systems. The paper closes by discussing the potential impact of the proposed initiatives.

2 Addressing Deep Learning Challenges through Integrations

Challenges for deep learning: Deep learning has achieved remarkable success in many task domains. In fact, at least under some conditions, deep learning can match or exceed human-level performance in face recognition [15], language translation [4], and game playing [14]. However, important challenges remain. LeCun pointed to three key challenges for deep learning:[2]

1. Learning with fewer labeled samples and/or fewer trials
2. Learning to reason
3. Learning to plan complex action sequences

Each of these is well-trodden ground for case-based reasoning. This suggests opportunities for integrations with case-based reasoning.

Models for Integrating Deep Learning with Symbolic Approaches: In his AAAI 2020 Engelmore lecture, Henry Kautz pointed to specific strengths of deep learning, such as learning hierarchically and that deep learning representations "directly support similarity." On the other hand, various other processes, such as combinatorial search, are natural for symbolic methods. In response to the divergent strengths, he advocated bringing together neural and symbolic traditions and proposed six possible combinations, including using the different technologies for specialized subroutines and a NeuroSymbolic approach in which symbolic rules are used to structure a neural system.[3] This paper proposes integrations at a more abstract level, in which the design of deep learning architectures is structured by the high-level CBR methodology.

---

[2] Quoted from
  https://drive.google.com/file/d/1r-mDL4IX_hzZLDBKp8_e8VZqD7fOzBkF/view
[3] https://www.cs.rochester.edu/u/kautz/talks/Kautz%20Engelmore%20Lecture.pdf

3 Implementing CBR with Deep Network Components

In the early days of case-based reasoning, CBR was often presented as an alternative technology to rule-based systems, associated with particular representation and implementation methods. In an influential paper, Ian Watson made a key observation: CBR can be implemented in many different ways using a range of methods. For example, CBR retrieval can be done using database technology [5]. Thus CBR is not a technology, but instead a *methodology:* a set of principles for a process of problem solving, interpretation, and learning that can be implemented using various technologies [16]. He frames the principles in terms of the classic four "REs" of the Aamodt and Plaza CBR cycle [1]: Retrieve similar cases, Reuse a similar case, Revise the solution to fit if necessary, and learn by Retaining. Each of these steps can be applied using multiple technologies.

Following this view, we can see CBR as a set of principles that could guide, for example, integrating multiple deep learning approaches to provide a CBR process for an end-to-end solution to a deep learning challenge problem. This shares aspects with multiple items in Kautz's categorization, but differs in that the defining aspect is not the specific technology, but rather the need for a particular functional sequence of processing steps.

The challenge for the CBR community is then to define the requisite tasks and integration. Various steps have been taken to bring deep learning components into CBR systems (e.g., [3, 8, 10, 13]). This challenge calls for an end-to-end effort to achieve CBR capabilities with a collection of deep learning-based components. This would have multiple benefits:

- **Providing CBR benefits while minimizing knowledge burdens:** CBR is no longer an alternative approach that loses the benefits of the knowledge-light processing of deep networks. When CBR is a unifying principle for guiding the design of deep learning systems, it can be implemented with the same technology.
- **Providing a framework for flexible technology integrations:** Even applying end-to-end CBR, the CBR process can still be implemented with whatever technology is most appropriate; the use of deep learning for some components does not preclude different technologies for others.
- **Providing increased flexibility through adaptation:** The *reasoning* part of CBR follows from case adaptation, the ability to transform solutions to new contexts. Explicitly integrating adaptation into deep learning systems could provide a new means for transfer.
- **Providing a new basis for learning from few examples:** Implementing a "true" case-based reasoning process able to reason and learn from single cases could help address the challenge of learning from limited data.
- **Providing a model for generating structured solutions:** Similarly, implementing a "true" case-based reasoning process able to manipulate structured cases could enable processing structured data such as action plans.
- **Reducing storage requirements**: Cases can capture knowledge compactly, in contrast to the potentially enormous requirements of networks.

If it is not possible to fully develop such a process within a deep learning architecture, hybrid solutions can still provide powerful processing capabilities.

4 Questions for a CBR-Based Pipeline

As discussed, the CBR methodology is agnostic to technology. However, applying that methodology in a neural network context requires addressing several key questions:

- **Case representation**: How can the rich structured cases of CBR be represented in a network context?
- **The role of cases**: What are the tradeoffs of explicit case retrieval rather than direct solution generation, and what are their respective roles? We consider this further below.
- **Case adaptation**: How and where should adaptation be applied? Can case adaptation be learned and applied within other processes of the CBR cycle, such as via adaptation at interior points of the network, rather than only to the retrieved solution? This might be seen as related to the question transformation of early CBR [6] and efforts at supporting analogical reasoning directly with embedded representations [11]. Neural networks have previously been applied to case adaptation [2], and recent efforts have applied deep learning to case adaptation using the case difference heuristic [8].
- **The meaning of the Retain step**: A fundamental principle of CBR is that results are retained as new cases. However, when learning is achieved by gradient-based training methods, the cost of learning by retraining after every case is prohibitive. Consequently, a core question is how to achieve lazy learning in a neural network context, or whether the case store must necessarily be implemented with another technology.

Developing a CBR-based pipeline raises the question of the role of explicit case representation and manipulation. Deep learning systems are eager learners; they receive (large quantities of) training data and learn weights that encode generalizations from that data. Case-based reasoners are lazy learners, retaining raw cases (or cases with limited processing) to re-use them. When using CBR to shape a deep learning pipeline, a natural question is the role of cases. There are three possibilities: To include an explicit case retrieval phase for "pre-packaged" cases; to "assemble" or generate cases by a reconstructive process without literal case storage (cf. [7]), or to dispense with explicit case retrieval/generation, solely adding an adaptation phase after solution generation, in the absence of cases. Adding adaptation to deep learning pipelines is an interesting — and potentially highly impactful — challenge for the CBR community. However, the full benefits of CBR, such as single-example lazy learning and explainability, require the use of cases.

5 CBR for AutoML

CBR can also be brought to deep learning—and other machine learning methods—in the context of automated machine learning (AutoML). AutoML focuses on methods to take as input a dataset, a challenge problem, and a library of primitives including machine learning algorithms, to automatically develop an end-to-end machine learning pipeline. It is being pursued by the DARPA Data-Driven Discovery of Models (D3M) program.[4] Most D3M teams focus on algorithm selection and hyperparameter optimization. However, some apply a "meta-learning" approach exploiting a database of prior solutions generated by the former "first principles" methods. This can be seen as a case base, and the CBR methodology, and specific lessons and methods for indexing, similarity, and adaptation, could play an important role in exploiting it. However, to our knowledge, this opportunity for synergy has not yet been pursued.

6 Conclusion: Future Paths

This challenge paper has proposed that beyond focusing on applying deep learning methods for CBR, the CBR community should focus on how the CBR methodology can help address the next generation of deep learning challenges. This may be especially beneficial for a CBR perspective on how to view problems and design architectures.

Bringing CBR to deep learning has the potential for great impact on future AI systems and to increase the reach of CBR. As a coarse-grained measure of the degree of attention to deep learning, a search of the Semantic Scholar archive of scholarly articles on May 6, 2020 yielded 11,500 results for "case-based reasoning" in the last five years, versus approximately 284,000 for "deep learning." Bringing case-based reasoning methodology to deep learning could also provide components for a new generation of knowledge-light CBR applications. Bringing CBR to AutoML provides an opportunity to harness strengths of CBR for effective use of multiple machine learning methods.

References

1. Aamodt, A., Plaza, E.: Case-based reasoning: Foundational issues, methodological variations, and system approaches. AI Communications 7(1), 39–52 (1994)
2. Corchado, J.M., Lees, B.: Adaptation of cases for case based forecasting with neural network support. In: Soft Computing in Case Based Reasoning, pp. 293–319. Springer (2001)

---

[4] https://www.darpa.mil/program/data-driven-discovery-of-models

3. Grace, K., Maher, M.L., Wilson, D.C., Najjar, N.: Combining CBR and deep learning to generate surprising recipe designs. In: Case-Based Reasoning Research and Development, ICCBR 2016. Springer, Berlin (2016)

4. Hassan, H., Aue, A., Chen, C., Chowdhary, V., Clark, J., Federmann, C., Huang, X., Junczys-Dowmunt, M., Lewis, W., Li, M., Liu, S., Liu, T.Y., Luo, R., Menezes, A., Qin, T., Seide, F., Tan, X., Tian, F., Wu, L., Wu, S., Xia, Y., Zhang, D., Zhang, Z., Zhou, M.: Achieving human parity on automatic chinese to english news translation (2018)

5. Kitano, H., Shimazu, H.: The experience sharing architecture: A case study in corporate-wide case-based software quality control. In: Leake, D. (ed.) Case-Based Reasoning: Experiences, Lessons, and Future Directions, pp. 235–268. AAAI Press, Menlo Park, CA (1996)

6. Kolodner, J.: Retrieval and Organizational Strategies in Conceptual Memory. Lawrence Erlbaum, Hillsdale, NJ (1984)

7. Leake, D.: Assembling latent cases from the web: A challenge problem for cognitive CBR. In: Proceedings of the ICCBR-11 Workshop on Human-Centered and Cognitive Approaches to CBR (2011)

8. Liao, C., Liu, A., Chao, Y.: A machine learning approach to case adaptation. In: 2018 IEEE First International Conference on Artificial Intelligence and Knowledge Engineering (AIKE). pp. 106–109 (2018)

9. López de Mántaras, R., McSherry, D., Bridge, D., Leake, D., Smyth, B., Craw, S., Faltings, B., Maher, M., Cox, M., Forbus, K., Keane, M., Aamodt, A., Watson, I.: Retrieval, reuse, revision, and retention in CBR. Knowledge Engineering Review 20(3) (2005)

10. Mathisen, B.M., Aamodt, A., Bach, K., Langseth, H.: Learning similarity measures from data. Progress in Artificial Intelligence (10 2019)

11. Mikolov, T., Chen, K., Corrado, G., Dean, J.: Efficient estimation of word representations in vector space. CoRR abs/1301.3781 (2013), `http://arxiv.org/abs/1301.3781`

12. Richter, M.: Introduction. In: Lenz, M., Bartsch-Spörl, B., Burkhard, H.D., Wess, S. (eds.) CBR Technology: From Foundations to Applications, chap. 1, pp. 1–15. Springer, Berlin (1998)

13. Sani, S., Wiratunga, N., Massie, S.: Learning deep features for knn-based human activity recognition. In: Proceedings of ICCBR 2017 Workshops (CAW, CBRDL, PO-CBR), Doctoral Consortium, and Competitions co-located with the 25th International Conference on Case-Based Reasoning (ICCBR 2017), Trondheim, Norway, June 26-28, 2017. CEUR Workshop Proceedings, vol. 2028, pp. 95–103. CEUR-WS.org (2017), `http://ceur-ws.org/Vol-2028/paper9.pdf`

14. Silver, D., Schrittwieser, J., Simonyan, K., Antonoglou, I., Huang, A., Guez, A., Hubert, T., Baker, L., atthew Lai, M., Bolton, A., Chen, Y., Lillicrap, T., Hui, F., Sifre, L., van den Driessche, G., Graepel, T., sabis, D.H.: Mastering the game of Go without human knowledge. Nature 550, 354–359 (19 October 2017)

15. Taigman, Y., Yang, M., Ranzato, M.A., Wolf, L.: Deepface: Closing the gap to human-level performance in face verification. In: IEEE Conference on Computer Vision and Pattern Recognition (2014)

16. Watson, I.: Case-based reasoning is a methodology not a technology. Knowledge-Based Systems 12(303-308) (1996)