DL-CBR Hybridization for Feature Generation and Similarity Assessment

Zachary Wilkerson

Indiana University, Bloomington, IN, 47408, United States

Abstract

Effective retrieval is essential to strong case-based reasoning performance, and retrieval quality is critically dependent on case indexing. Such indices are not always feasible to generate manually, and so a thorough exploration of how features and weights may be generated automatically (especially using deep learning) is necessary. To that end, this summary outlines a research plan for investigating structural influences on feature quality, how learned features may be used in concert with knowledge-engineered features, and how weights may be generated in feature-dense spaces created by feature learning. It also proposes a methodology for modular exploration of various models, training set sizes, numbers of features generated, etc., to provide a comprehensive foundation of index generation using deep learning. Finally, it points to already-published research that works towards some of these goals and illustrates how results from these existing projects inform future research plans.

Keywords

Case-based reasoning, Deep learning, Hybrid systems, Feature learning, Weight learning

1. Introduction

Case-based reasoning (CBR) performance relies significantly on retrieving useful cases from the case base. In turn, retrieval quality depends on indices used to characterize/discriminate between cases. High-quality indices can be derived through manual knowledge engineering (e.g., [1, 2]), but this approach can be costly and is not feasible in some domains. For example, indexing vocabularies may be unsatisfactory for poorly-understood domains or for complex tasks such as computer vision. An analogous challenge exists for weights as well–effective feature weights can augment feature information for indexing, but even provided a comprehensive feature set, it can be difficult and/or expensive to identify useful weighting information for those features.

These problems may be addressed using feature and weight learning. Initially, this was achieved using symbolic methods (e.g., [3]); however, with recent advances in deep learning (DL, esp. in domains such as computer vision), it is natural to consider how increased performance of DL architectures may be translated to CBR retrieval. Specifically, CBR systems can be described as inherently interpretable via case presentation, but black-box DL systems are traditionally viewed as more accurate for most domains; however, if a CBR system can be made more accurate by leveraging DL methods/structures, then resulting DL-CBR hybrid systems may be applicable to a wide variety of domains, representing a "best of both worlds" with regards

ICCBR DC'22: Doctoral Consortium at ICCBR-2022, September, 2022, Nancy, France
zachwilk@indiana.edu (Z. Wilkerson)

to accuracy and explainability (e.g., "Twin systems" by Kenny and Keane in [4]). In this vein, recent research leverages neural networks to generate and refine feature information inferred from training examples for classification and/or explanation [5, 6, 7], but there still exist many areas of potential DL-CBR integration for which there has been little or no research exploration.

This research summary outlines strategies for exploring feature and weight learning using DL in greater depth, presenting both a blueprint detailing potential research objectives and methodologies as well as an overview of steps taken so far and the resulting publications.

2. Research Objectives and Methodology

This research broadly seeks to deeply explore methods for leveraging DL models to generate indexing information to supplement or replace information gathered through knowledge engineering for computer vision-related tasks. This overarching goal can be subdivided into the four primary investigation regions described below, which are followed by a proposed research plan outline for exploring them.

2.1. Exploring Methods for Generating High-Quality Features

Under the umbrella of network-based feature generation, there exist multiple potential variables that can influence feature quality. For one, different model structures provide unique pathways for feature generation, exemplified by different computer vision approaches (e.g., comparing AlexNet, Inception, DenseNet, transformers, and MLP-based models); additionally, architectures may be leveraged in different ways (e.g., using ensembles of networks for localized feature generation in the multi-net approach explored in [8]). For another, the way in which features are extracted from DL models affects feature quality (e.g., extracting features from different locations in a model). Analysis of the impact of such variables on feature quality is an essential foundation for optimizing DL-CBR hybrid system performance.

2.2. Using Knowledge-Engineered and Learned Features in Concert

While feature learning can be useful in domains for which generating features through knowledge-engineering is not feasible, more research is required to evaluate feature learning augmenting incomplete knowledge-engineered indices. This includes exploring methods for effectively using both feature sets in concert, for which it may be necessary to mitigate harmful effects of a "curse of dimensionality" as a result of extracted feature spaces being generally denser than knowledge-engineered ones. Integrating knowledge-engineered and network-generated features also requires investigation into potential discretization of continuous network-generated features as well as into the independence of generated features and/or their correlation with knowledge-engineered features.

2.3. Refining Weight Learning Methods for Feature-Dense Spaces

As mentioned above, using learned features can result in similarity assessment being performed in feature-dense spaces, for which conventional methods of weight learning (e.g., [4]) may be less effective. To this end, it is important to consider ways in which such techniques may be refined to accommodate larger numbers of features, and/or to explore methods to extract weights directly from a network architecture, potentially in concert with feature extraction. This objective also encompasses how combinations of feature and weight learning methods influence retrieval quality, especially with regards to extracting feature weights from different network architectures.

2.4. Evaluating DL-CBR Hybrid Model Explainability

Investigation of hybrid systems leveraging interpretable CBR structures alongside more opaque DL systems demands contextualization relative to explainability. Innately, DL-CBR hybrid systems imply an overall architecture that is more interpretable than an out-of-the-box DL model but less so than a CBR system using only knowledge-engineered information. Thus, in addition to optimizing index quality to maximize retrieval accuracy, it is important to assess where on an explainability spectrum that this work sits and to take measures where possible to maximize interpretability.

2.5. Proposed Research Plan

The four research objectives described above encompass a diverse range of research avenues, and they present specific sub-goals that align with an overarching three-step process that guides the proposed research methodology. Specifically, this process begins with deep exploration of network-based index generation methods, including different ways in which generated indices may be integrated into a CBR system. The second step involves post-processing of these indices, particularly for optimizing CBR system performance/accuracy, but also potentially including discretization for better combination with knowledge-engineered index information where applicable. Finally, the resulting DL-CBR hybrid model is analyzed/contextualized with respect to explainability, especially in comparison to CBR systems using only knowledge-engineered indexing information.

To this end, initial experiments imitating established index generation methods (e.g., [5, 7], see next section for details) have established both a proof of concept that network-generated and knowledge-engineered features used in concert can enable greater CBR accuracy than either feature set used individually and that the network architecture/structure can have a substantial impact on feature quality. Next steps will focus on other network structure influences (esp. how different DL models affect feature quality), enabling a comprehensive analysis of network-generated feature sets augmenting knowledge-engineered feature sets.

Beyond this point, future experiments could continue in any of several directions. For one, feature weighting strategies may be more deeply investigated and/or revised in the context of potentially denser feature spaces created using network-generated features. For another, potential relationships between knowledge-engineered and network-generated features may be explored for explainability purposes. Such investigations would include dependency correlations between knowledge-engineered and network-generated features and/or methods by which continuous network-generated features may be discretized, along with the resulting impact on feature quality.

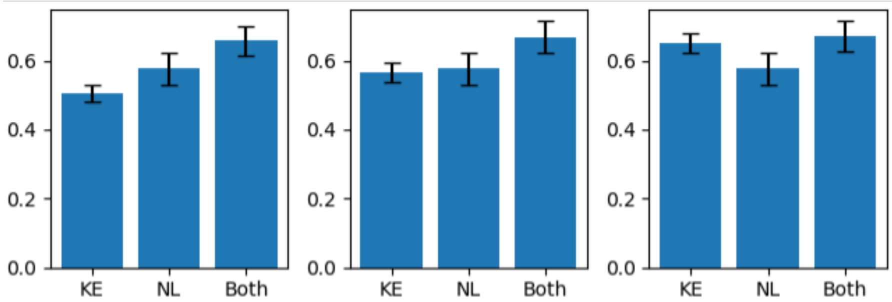

Figure 1: Comparison of retrieval accuracy using knowledge-engineered (KE) features, network-generated/learned (NL) features, or both sets together (First published in Case-Based Reasoning Research and Development, ICCBR 2021 by Springer [9]). From right to left, knowledge-engineered features are perturbed by greater magnitudes, resulting in less per-feature reliability.

3. Progress to Date

As of the writing of this summary, most progress to date has focused on providing a conceptual foundation and initial empirical tests for DL-CBR hybridization for retrieval, in line with the objectives presented above. Specifically, such explorations include using network-generated features in concert with knowledge-engineered features for greater classification accuracy, as well as investigating structural influences (e.g., network architecture/structure and feature extraction location) on feature quality, both using retrieval accuracy as proxy. The following subsections summarize the associated publications.

3.1. Augmenting Similarity Feature Engineering with Deep Learning

This research [9] assumes the availability of knowledge-engineered feature information for a given domain, but that such feature information may be incomplete and/or inaccurate. In these instances, existing feature information may be supplemented by additional learned features extracted from raw data using DL. In these instances, the inclusion of learned features improves retrieval accuracy by capturing indexing information to which humans might not be sensitive.

Results supporting this hypothesis are obtained using a zero-shot learning dataset for computer vision. Each image is associated with a unique case, and per-class feature information from the dataset is perturbed based on a random coefficient and combined with values extracted from the image using a convolutional neural network (CNN) to form the case's feature set. The two combined sets of values represent knowledge-engineered and network-generated features, respectively. Retrieval accuracy values for the aggregated feature set are compared against corresponding accuracy values using either component set exclusively (Figure 1).

Based on the outcomes from these initial tests, combining feature sets does improve retrieval accuracy. However, additional variables such as the reliability of knowledge-engineered features and the number of features extracted from the CNN may significantly influence the magnitude of accuracy improvement. In addition, preliminary tests regarding weight extraction in parallel with feature extraction suggest that more feature-dense spaces created by extracting features

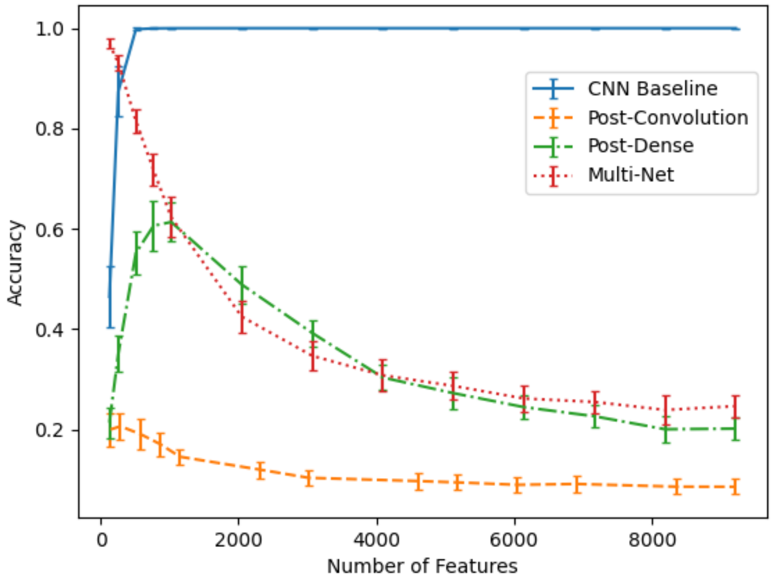

Figure 2: Comparison of different feature extraction methods with an end-to-end CNN classifier for different numbers of features for a basic CNN architecture (First published in Case-Based Reasoning Research and Development, ICCBR 2022 by Springer Nature [8]). Post-Convolution extracts features immediately after convolution and pooling steps, Post-Dense extracts features immediately after the dense layers (before the output layer), and Multi-Net extracts features as in Post-Dense from an ensemble of networks based on a candidate retrieved case's class.

from CNNs seem to accommodate established weighting strategies poorly.

3.2. Exploring Structural Influences on Generated Feature Effectiveness

In contrast to the previous description, this work [8] specifically investigates how the way in which features are extracted affects feature quality. To this end, two feature extraction locations and a novel model ensemble structure that generates localized features are explored.

The experiments investigate the hypothesis that extracting features from later in the network results in higher-quality features from the perspective of the CBR system. Retrieval accuracy values are used as proxy for feature quality among the three proposed methods and a CNN baseline (Figure 2). Additionally, keeping in mind the consequences of feature-dense spaces discovered in the previous work, varying numbers of features are extracted from each model.

The results support the hypothesis and suggest that localized feature sets may be especially accurate for feature-sparse scenarios preferred by CBR systems (if at the cost of increased training time). The number of features does significantly impact retrieval performance as well, both as a "curse of dimensionality" for large numbers of features and as a minimum requirement for DL model convergence for smaller numbers of features.

4. Future Work

Future work will build on the current publications' findings while moving forward to address other objectives. In the short term, research will focus on exploring feature generation using both different DL models across multiple datasets and different experimental parameters (e.g., number of training examples). Later experiments will investigate potential weight generation methods, as well as how weight generation methods are affected by the number of features.

References

[1] D. Leake, An indexing vocabulary for case-based explanation, in: Proceedings of the Ninth National Conference on Artificial Intelligence, AAAI Press, Menlo Park, CA, 1991, pp. 10–15.

[2] R. Schank, M. Brand, R. Burke, E. Domeshek, D. Edelson, W. Ferguson, M. Freed, M. Jona, B. Krulwich, E. Ohmayo, R. Osgood, L. Pryor, Towards a general content theory of indices, in: Proceedings of the 1990 AAAI spring symposium on Case-Based Reasoning, AAAI Press, Menlo Park, CA, 1990.

[3] S. Bhatta, A. Goel, Model-based learning of structural indices to design cases, in: Proceedings of the IJCAI-93 Workshop on Reuse of Design, IJCAI, Chambery, France, 1993, pp. A1–A13.

[4] E. M. Kenny, M. T. Keane, Twin-systems to explain artificial neural networks using case-based reasoning: Comparative tests of feature-weighting methods in ANN-CBR twins for XAI, in: Proceedings of the Twenty-Eighth International Joint Conference on Artificial Intelligence, 2019.

[5] S. Sani, N. Wiratunga, S. Massie, 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, volume 2028 of *CEUR Workshop Proceedings*, CEUR-WS.org, 2017, pp. 95–103.

[6] K. Martin, N. Wiratunga, S. Sani, S. Massie, J. Clos, A convolutional siamese network for developing similarity knowledge in the selfBACK dataset., in: A. A. Sanchez-Ruiz, A. Kofod-Petersen (Eds.), Proceedings of the ICCBR 2017 Workshop on Case-Based Reasoning and Deep Learning, CEUR Workshop Proceedings, 2017, pp. 85–94.

[7] J. T. Turner, M. W. Floyd, K. M. Gupta, T. Oates, NOD-CC: A hybrid CBR-CNN architecture for novel object discovery, in: Case-Based Reasoning Research and Development, ICCBR 2019, 2019, pp. 373–387.

[8] D. Leake, Z. Wilkerson, D. Crandall, Extracting case indices from convolutional neural networks: A comparative study, in: Case-Based Reasoning Research and Development, ICCBR 2022 (in print), 2022.

[9] Z. Wilkerson, D. Leake, D. Crandall, On combining knowledge-engineered and network-extracted features for retrieval, in: Case-Based Reasoning Research and Development, ICCBR 2021, 2021, pp. 248–262.