Categories of Differentiable Polynomial Circuits for Machine Learning

Paul Wilson[1,2][0000−0003−3575−135X] and Fabio Zanasi[2][0000−0001−6457−1345]

[1] University of Southampton
[2] University College London

Abstract. Reverse derivative categories (RDCs) have recently been shown to be a suitable semantic framework for studying machine learning algorithms. Whereas emphasis has been put on training methodologies, less attention has been devoted to particular *model classes*: the concrete categories whose morphisms represent machine learning models. In this paper we study presentations by generators and equations of classes of RDCs. In particular, we propose *polynomial circuits* as a suitable machine learning model. We give an axiomatisation for these circuits and prove a functional completeness result. Finally, we discuss the use of polynomial circuits over specific semirings to perform machine learning with discrete values.

1 Introduction

Reverse Derivative Categories [7] have recently been introduced as a formalism to study abstractly the concept of differentiable functions. As explored in [8], it turns out that this framework is suitable to give a categorical semantics for gradient-based learning. In this approach, models–as for instance neural networks–correspond to morphisms in some RDC. We think of the particular RDC as a 'model class'–the space of all possible definable models.

However, much less attention has been directed to actually defining the RDCs in which models are specified: existing approaches assume there is some chosen RDC and morphism, treating both essentially as a black box. In this paper, we focus on classes of RDCs which we call 'polynomial circuits', which may be thought of as a more expressive version of the boolean circuits of Lafont [9], with wires carrying values from an arbitrary semiring instead of $\mathbb{Z}_2$. Because we ensure polynomial circuits have RDC structure, they are suitable as machine learning models, as we discuss in the second part of the paper.

Our main contribution is to provide an algebraic description of polynomial circuits and their reverse derivative structure. More specifically, we build a presentation of these categories by operation and equations. Our approach will proceed in steps, by gradually enriching the algebraic structures considered, and culminate in showing that a certain presentation is *functionally complete* for the class of functions that these circuits are meant to represent.

An important feature of our categories of circuits is that morphisms are specified in the graphical formalism of *string diagrams*. This approach has the benefit of making the model specification reflect its combinatorial structure. Moreover, at a computational level, the use of string diagrams makes available the principled mathematical toolbox of *double-pushout rewriting*, via an interpretation of string diagrams as hypergraphs [3,4,5]. Finally, the string diagrammatic presentation suggests a way to encode polynomial circuits into datastructures: an important requirement for being able to incorporate these models into tools analogous to existing deep learning frameworks such as TensorFlow [1] and PyTorch [10].

Tool-building is not the only application of the model classes we define here. Recent neural networks literature [2,6] proposes to improve model performance (e.g. memory requirements, power consumption, and inference time) by 'quantizing' network parameters. One categorical approach in this area is [11], in which the authors define learning directly over boolean circuit models instead of training with real-valued parameters and then quantizing. The categories in our paper can be thought of as a generalisation of this approach to arbitrary semirings.

This generalisation further yields another benefit: while neural networks literature focuses on finding particular 'architectures' (i.e. specific morphisms) that work well for a given problem, our approach suggests a new avenue for model design: changing the underlying semiring (and thus the corresponding notion of arithmetic). To this end, we conclude our paper with some examples of finite semirings which may yield new approaches to model design.

2 Full Paper

The details of our approach are to appear in a full paper accepted for publication in the proceedings of ICGT 2022 (see https://icgt2022.gitlab.io/). A preprint is available on the Arxiv at https://arxiv.org/abs/2203.06430

References

1. Abadi, M., et al.: TensorFlow: Large-scale machine learning on heterogeneous systems (2015), https://www.tensorflow.org/
2. Bengio, Y., Léonard, N., Courville, A.: Estimating or propagating gradients through stochastic neurons for conditional computation (2013). https://doi.org/10.48550/ARXIV.1308.3432, https://arxiv.org/abs/1308.3432
3. Bonchi, F., Gadducci, F., Kissinger, A., Sobocinski, P., Zanasi, F.: String diagram rewrite theory i: Rewriting with frobenius structure (2020). https://doi.org/10.48550/ARXIV.2012.01847
4. Bonchi, F., Gadducci, F., Kissinger, A., Sobocinski, P., Zanasi, F.: String diagram rewrite theory ii: Rewriting with symmetric monoidal structure (2021). https://doi.org/10.48550/ARXIV.2104.14686
5. Bonchi, F., Gadducci, F., Kissinger, A., Sobociski, P., Zanasi, F.: String diagram rewrite theory iii: Confluence with and without frobenius (2021). https://doi.org/10.48550/ARXIV.2109.06049

6. Choi, J., et al.: Accurate and efficient 2-bit quantized neural networks. In: Talwalkar, A., Smith, V., Zaharia, M. (eds.) Proceedings of Machine Learning and Systems. vol. 1, pp. 348–359 (2019), https://proceedings.mlsys.org/paper/2019/file/006f52e9102a8d3be2fe5614f42ba989-Paper.pdf
7. Cockett, R., Cruttwell, G., Gallagher, J., Lemay, J.S.P., MacAdam, B., Plotkin, G., Pronk, D.: Reverse derivative categories (2019). https://doi.org/10.48550/ARXIV.1910.07065
8. Cruttwell, G.S.H., Gavranovi, B., Ghani, N., Wilson, P., Zanasi, F.: Categorical foundations of gradient-based learning (2021). https://doi.org/10.48550/ARXIV.2103.01931, https://arxiv.org/abs/2103.01931
9. Lafont, Y.: Towards an algebraic theory of Boolean circuits. Journal of Pure and Applied Algebra **184**(2-3), 257–310 (Nov 2003). https://doi.org/10.1016/S0022-4049(03)00069-0
10. Paszke, A., et al.: Pytorch: An imperative style, high-performance deep learning library. In: Advances in Neural Information Processing Systems 32, pp. 8024–8035. Curran Associates, Inc. (2019)
11. Wilson, P., Zanasi, F.: Reverse derivative ascent: A categorical approach to learning boolean circuits. Electronic Proceedings in Theoretical Computer Science **333**, 247260 (Feb 2021). https://doi.org/10.4204/eptcs.333.17