Recognizing algebraic properties from multiplication tables

Sujoy Mukherjee[1], Daniel Scofield[2], and Petr Vojtěchovský[3]*

[1] University of Denver, Denver, Colorado, U.S.A.
sujoy.mukherjee@du.edu
[2] Francis Marion University, Florence, South Carolina, U.S.A.
daniel.scofield@fmarion.edu
[3] University of Denver, Denver, Colorado, U.S.A.
petr.vojtechovsky@du.edu

Abstract

We study the problem of recognizing algebraic properties (such as commutativity, associativity, latin square property and self-distributivity) from multiplication tables by means of neural networks. We achieve reasonable accuracy but the problem is highly sensitive to the structure of training data. Our analysis sheds some light on the question of what exactly are neural networks learning and on the problem of how to make neural networks learn what we want them to learn (rather than a proxy property). This is work in progress—more details will be reported at the time of the conference.

1 General introduction

Research into applications of machine learning and artificial intelligence has increased dramatically in the past ten years, enabled by the availability of computing power and the accessibility of software packages, etc. Many attempts have been made to solve problems in pure mathematics using these methods [5]. Davies et al. proposed a general framework for guiding human math intuition using AI [1].

It has been observed that machine learning models may achieve high levels of accuracy by exploiting spurious correlations or artifacts in training data [4], or may discover "shortcut" rules that succeed within the realm of experimental test/train data but fail to generalize to simple examples from a slightly different distribution [2]. Thus caution is warranted when considering results from ML experiments.

2 Recognizing algebraic properties

We are interested in using ML to recognize certain algebraic properties from multiplication tables.

A given algebraic property P typically implies many other properties which, unlike P, might be easy to glean from simple statistical tests. (For instance, groups contain an identity element and their multiplication tables are latin squares.) If the implied properties are not carefully controlled for during training and validation, it seems likely that ML will learn to recognize some of the consequences of P rather than P itself.

He and Kim [3] report that ML can classify with high degree of accuracy whether multiplication tables belong to certain finite groups. In particular, they deduce that ML can learn to recognize associativity.

*Supported by the Simons Foundation Mathematics and Physical Sciences Collaboration Grant for Mathematicians no. 855097

Being somewhat pessimistic of such claims, we set out to investigate the situation more systematically on small multiplication tables. Here is a sampling of our observations for the case of latin squares:

- As a base case, given two disjoint sets A and B of random multiplication tables (of a fixed size), ML does not seem to be able to learn membership in A.
- Training ML on multiplication tables that are either latin squares or are far from being latin squares results in a fairly accurate recognition (exceeding 95 percent) of the latin property.
- But the model trained as above has a high rate of false positives when tested against multiplication tables that are close to being latin squares (which we can think of as hard counterexamples).
- The rate of false positives improves dramatically when the training data consists only of latin squares and multiplication tables that are close to being latin squares. Interestingly, models so trained continue to perform very well on recognizing latin squares among general multiplication tables, despite never being exposed to general multiplication tables during training.

This suggests a general strategy for the recognition of a property P by ML, in which the training data consists of multiplication tables satisfying P and of multiplication tables that are close (in some sense) to satisfying P. We will report on the results for various choices of P, such as commutativity, associativity and self-distributivity. We will also comment on some previously used training techniques (such as permuting the rows, columns and symbols of multiplication tables by three independent permutations) which are not necessarily mathematically sound, sometimes in a subtle way.

References

[1] Davies, A., Veličković, P., Buesing, L. et al., *Advancing mathematics by guiding human intuition with AI*, Nature **600**, 70–74 (2021).

[2] Geirhos, R., Jacobsen, JH., Michaelis, C. et al. *Shortcut learning in deep neural networks*, Nat Mach Intell **2**, 665–673 (2020).

[3] Yang-Hui He and Minhyong Kim, *Learning algebraic structures: Preliminary investigations*, International Journal of Data Science in the Mathematical Sciences **01**, no. **01**, 3–22 (2023)

[4] Lapuschkin, S., Wäldchen, S., Binder, A. et al., *Unmasking Clever Hans predictors and assessing what machines really learn*, Nat Commun **10**, 1096 (2019).

[5] Williamson, G., *Is deep learning a useful tool for the pure mathematician?*, Bulletin of the American Mathematical Society, Volume **61**, Number **2**, April 2024, Pages 271–286.