

ACCEPTED MANUSCRIPT • OPEN ACCESS

Transforming the bootstrap: using Transformers to compute scattering amplitudes in planar $N = 4$ Super Yang-Mills theory

To cite this article before publication: Tianji Cai *et al* 2024 *Mach. Learn.: Sci. Technol.* in press <https://doi.org/10.1088/2632-2153/ad743e>

Manuscript version: Accepted Manuscript

Accepted Manuscript is “the version of the article accepted for publication including all changes made as a result of the peer review process, and which may also include the addition to the article by IOP Publishing of a header, an article ID, a cover sheet and/or an ‘Accepted Manuscript’ watermark, but excluding any other editing, typesetting or other changes made by IOP Publishing and/or its licensors”

This Accepted Manuscript is © 2024 The Author(s). Published by IOP Publishing Ltd.



As the Version of Record of this article is going to be / has been published on a gold open access basis under a CC BY 4.0 licence, this Accepted Manuscript is available for reuse under a CC BY 4.0 licence immediately.

Everyone is permitted to use all or part of the original content in this article, provided that they adhere to all the terms of the licence <https://creativecommons.org/licenses/by/4.0>

Although reasonable endeavours have been taken to obtain all necessary permissions from third parties to include their copyrighted content within this article, their full citation and copyright line may not be present in this Accepted Manuscript version. Before using any content from this article, please refer to the Version of Record on IOPscience once published for full citation and copyright details, as permissions may be required. All third party content is fully copyright protected and is not published on a gold open access basis under a CC BY licence, unless that is specifically stated in the figure caption in the Version of Record.

View the [article online](#) for updates and enhancements.

Transforming the Bootstrap: Using Transformers to Compute Scattering Amplitudes in Planar $\mathcal{N} = 4$ Super Yang-Mills Theory

Tianji Cai^{a*}†, Garrett W. Merz^{b*}†, François Charton^{c*}, Niklas Nolte^c, Matthias Wilhelm^d, Kyle Cranmer^b, Lance J. Dixon^a

^a SLAC National Accelerator Laboratory

^b Data Science Institute, University of Wisconsin-Madison

^c FAIR, Meta

^d Niels Bohr Institute, University of Copenhagen

tianji@slac.stanford.edu, garrett.merz@wisc.edu, fcharton@meta.com, nolte@meta.com, matthias.wilhelm@nbi.ku.dk, kyle.cranmer@wisc.edu, lance@slac.stanford.edu

May 2024

Abstract. We pursue the use of deep learning methods to improve state-of-the-art computations in theoretical high-energy physics. Planar $\mathcal{N} = 4$ Super Yang-Mills theory is a close cousin to the theory that describes Higgs boson production at the Large Hadron Collider; its scattering amplitudes are large mathematical expressions containing integer coefficients. In this paper, we apply Transformers to predict these coefficients. The problem can be formulated in a language-like representation amenable to standard cross-entropy training objectives. We design two related experiments and show that the model achieves high accuracy ($> 98\%$) on both tasks. Our work shows that Transformers can be applied successfully to problems in theoretical physics that require exact solutions.

1. Introduction

Particle physics at the energy frontier is entering an exciting new era of high-precision experiments, ushered in by the high-luminosity upgrade of the Large Hadron Collider (LHC). Exploiting the full physics potential of the experimental data requires substantial improvements in the predictions of Standard Model (SM) [1] processes, both as backgrounds to new physics, and for measuring Higgs boson couplings and other SM parameters.

Many ingredients are necessary for these predictions, see e.g. [2] for a review. At the heart of all such calculations are scattering amplitudes—the fundamental quantum-mechanical building blocks for transition probabilities between asymptotic states. The

*Denotes equal contribution

†Authors to whom correspondence should be addressed

Transforming the Bootstrap

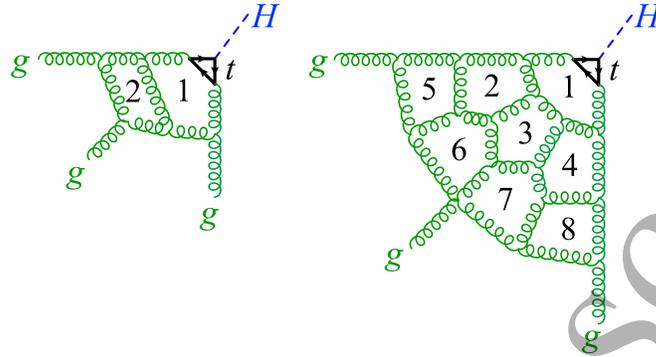


Figure 1. Sample Feynman diagrams for the process $gg \rightarrow Hg$ at two loops (left) and eight loops (right) in QCD. The same diagrams contribute in SYM, where the Higgs boson H and top quark (t) triangle is replaced by a particular local operator in the theory, and the process is referred to as a form factor.

conventional way to compute scattering amplitudes uses Feynman diagrams (see Figure 1 for examples), which graphically organize a series of terms in a perturbative expansion. Performing high-precision calculations in the theory of quantum chromodynamics (QCD) requires Feynman diagrams containing at least two loops [3–10]. Each loop represents intermediate-state virtual particles whose unobserved momenta must be integrated over. Each successive order of precision demands the addition of another loop to the diagram. Unfortunately, the number of possible Feynman diagrams, and thus the number of integrals that must be performed, grows factorially with loop order, quickly making these calculations intractable.

A recently-developed alternative technique, known as the *amplitude bootstrap* [11–13], attempts to directly construct candidate solutions for multi-loop amplitudes. It has mainly been applied so far to a simpler relative of QCD, called *planar $\mathcal{N} = 4$ super-Yang-Mills theory* (SYM). The amplitude bootstrap circumvents many of the computational and numerical challenges that arise from the Feynman diagram approach. It leverages the rich, yet highly constrained, analytical structure of amplitudes that arises from the particular recurrent features of the integrals involved. Using this technique, the form of the amplitude can be determined *a priori*, and the finite-dimensional solution space at a given loop order can be strongly constrained through a large system of linear relations with integer coefficients. Many of the linear relations are found by analyzing the lower-loop results. A small number of additional constraints, derived from behavior in physical limits, can then be applied in order to obtain a unique solution.

The amplitude bootstrap allows for the computation of amplitudes up to eight loops in SYM [14], *vs.* two loops using traditional Feynman diagram methods for the same quantity in QCD [15], as depicted in Figure 1. However, when using the bootstrap technique, the number of linear relations to be solved and the number of unknown coefficients both increase by a factor of about 4 at each subsequent loop order [14]. Since

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

Transforming the Bootstrap

3

it is not possible to determine the minimal set of independent equations in advance, the number of relations can be several times larger than the number of unknown coefficients. The computational cost of generating and solving the equations thus increases by at least a factor of $4^2 = 16$ at each loop order. Performing computations by this method becomes infeasible beyond about loop $L = 9$. This situation necessitates the development of new methods that may exploit hitherto unobserved patterns in the data in order to simplify the computation of scattering amplitudes at higher loops.

Notably, amplitudes in SYM can be expressed as sets of tokenizable “words” with integer coefficients; see Section 2. We refer to these words as *keys* which index into the coefficients, and to a key-coefficient pair as an *element*. Exploiting the many linear relations between these elements amounts to solving a large system of linear equations (as explained above), where the integer solutions are hard to discover, but easy to verify.

In our case, verification includes first checking that the proposed expression is a legitimate function in the appropriate space (see Appendix A), and then comparing its behavior in a physical limit to the predictions of the form-factor operator product expansion (FFOPE) [16, 17]. The latter constraint is stringent enough to fix the answer uniquely [14].

Since both the keys and the integers can be represented as sequences of tokens, we can train deep learning models such as Transformers to predict the coefficient associated with each key.

Transformers [18] are incredibly versatile neural network architectures that employ an *attention mechanism* [19] to learn complex nonlinear relationships between input features. They have revolutionized fields ranging from natural language processing [20] and computer vision [21] to formal symbolic mathematics [22, 23]. Inspired by these many recent successes, we apply Transformers to two sets of experiments; see Section 3 for our model setup.

We first show in Section 4 that Transformers can accurately predict elements of the solution at a given loop order when trained on other elements at the same loop. In Section 5 we show that the model is able to do this even when data is presented in a highly compressed format. To make these predictions successfully, many features of the complex relationships between individual terms must be learned by the model. In the subsequent Section 6, we explore how a number of the known linear relations are learned as a function of training epoch, and use this information to draw conclusions about the learning dynamics of the model. In Section 7 we show that augmenting a small amount of training data at a given loop with data from a lower loop improves performance, and discuss prospects for future multi-loop experiments.

The mathematical structure of the problem hints that some elements of the solution at higher loops may be determined using related elements of the solution at lower loops. Our second goal is therefore to discover this relationship implicitly, assuming it exists. In Section 8 we train Transformers to predict coefficients of terms at loop $L + 1$, given a set of coefficients of potentially related terms at loop L . We also perform a number of ablations to determine conditions under which this relationship is no longer learnable.

Transforming the Bootstrap

This study allows us to uncover certain features of the cross-loop relationship, which may prove crucial in further developing the bootstrap program.

We emphasize that the goal of these experiments is not to optimize performance. Rather it is to show that properties of amplitudes can be learned by AI models, and to probe the learning dynamics.

Additionally, this paper contains several Appendices. In [Appendix A](#), we further describe the mathematical formalism by which scattering amplitudes are expressible in a language-like fashion, via their *symbols*, which are sums of pairs of keys and integer coefficients. In [Appendix B](#), we give a more comprehensive list of linear relations between symbol terms and evaluate them as a function of training epoch. In [Appendix C](#) we perform a number of architecture ablations and evaluate their effects on the coefficient prediction task of Section [4](#). In [Appendix D](#) we perform additional ablation experiments in the manner of Section [8](#) in order to further characterize the correspondence between elements at different loop orders.

1.1. Related Work

The amplitude bootstrap program has a long tradition, dating back to ref. [\[11\]](#). A recent review can be found in ref. [\[24\]](#), and the specific data we use is from ref. [\[14\]](#). Our work supplements the traditional approach by offering a novel problem-solving framework, where human intelligence is augmented by artificial intelligence to further push the state-of-the-art for amplitude calculations.

In a similar spirit, a number of recent works have also leveraged deep learning to tackle *analytical* calculations in theoretical physics. In particular, a sequence-to-sequence Transformer has been used to compute the squared amplitude of a particle interaction symbolically [\[25\]](#); deep reinforcement learning has been applied to explore the landscape of string vacua [\[26\]](#); and Transformers have been employed to simplify polylogarithms [\[27\]](#), which are complicated mathematical functions entering multi-loop amplitudes similar to those we study (see Section [2.1](#)). However, no previous work has used Transformers to perform computations in the amplitude bootstrap paradigm.

Methodologically, our work is closely related to recent works using Transformers for *symbolic mathematical data*. For example, Transformers have been taught to perform mathematical tasks such as solving differential equations [\[28\]](#), learning recurrent sequences [\[29\]](#), and finding the greatest common divisor of number pairs [\[30\]](#). A comparable approach has also been used to solve linear algebra tasks [\[23\]](#), including eigenvector decomposition and matrix inversion, which share many structural similarities with our amplitude bootstrap method. Additionally, our first experiment, in which we use some elements of a scattering amplitude to predict others, can be framed as a tensor completion problem—at a fixed loop order, the model must learn to fill in unseen elements of the solution based on elements it has seen at training time. This task is similar in some respects to low-rank matrix completion [\[31\]](#), where Transformers have previously been employed [\[32\]](#).

2. Three-Gluon Form Factors in Planar $\mathcal{N} = 4$ SYM Theory

The amplitude bootstrap program has seen substantial success in planar $\mathcal{N} = 4$ super-Yang-Mills theory [33]. Similar to QCD, SYM contains gluons which self-interact; but instead of including quarks, it contains four gluinos and six scalars. The gluons, gluinos and scalars are all massless, and all transform into each other under the $\mathcal{N} = 4$ supersymmetry. They all have the same number of internal “color” degrees of freedom. We take the number of colors N_c to infinity, and refer to the Feynman diagrams that contribute to the scattering of these massless particles as *planar*. As a theoretical laboratory or model system for QCD, SYM allows us to see much further into the perturbative expansion than QCD. For example, a class of SYM amplitudes was recently computed to *eight* loops [14, 34]. These amplitudes, referred to as *three-gluon form factors* $\mathcal{F}_{3\text{gFF}}$, involve three massless gluons and a massive color-singlet operator. The operator couples to gluons very similarly to how the Higgs boson does in the limit of a very heavy top quark. Thus $\mathcal{F}_{3\text{gFF}}$ is the SYM analog of the QCD process $gg \rightarrow Hg$, which is known only to two loops [15]. In fact, part of the QCD form-factor result (the so-called “highest-weight” part) is identical to the SYM result [35, 36]. Figure 1 shows sample Feynman diagrams for this process to the current highest calculable loop order in QCD and SYM, respectively.

In this work, we focus exclusively on the three-gluon form factors $\mathcal{F}_{3\text{gFF}}$ in planar $\mathcal{N} = 4$ SYM. The known results up to loop $L = 7$ are used for model training and evaluation.

2.1. Symbols: A Simple Language for Amplitudes

The three-gluon form factors, like many other amplitudes in SYM, can be expressed in terms of functions called *generalized polylogarithms*. They are multiple, iterated integrations of rational functions. In SYM, the calculation of scattering amplitudes at loop order L requires integrals that are iterated $2L$ times. Such amplitudes $\mathcal{F}^{(L)}$ can be characterized by another mathematical object known as the *symbol* [37]:

$$\mathcal{S}[\mathcal{F}^{(L)}] = \sum_{l_{i_1}, \dots, l_{i_{2L}} \in \mathcal{L}_m} C^{l_{i_1}, \dots, l_{i_{2L}}} l_{i_1} \otimes \dots \otimes l_{i_{2L}}. \quad (1)$$

Here $\mathcal{L}_m = \{l_1, \dots, l_m\}$ is the *symbol alphabet* containing m letters l_i , which are in turn functions of the particles’ four-momenta, and $C^{l_{i_1}, \dots, l_{i_{2L}}}$ is a $2L$ -fold tensor of integer coefficients, most of which are zero. In other words, a solution for the symbol at loop L can be represented by m^{2L} integers, with each sequence of $2L$ letters (i.e., $l_{i_1} \otimes \dots \otimes l_{i_{2L}}$) serving as a key indexing into the integer-valued tensor $C^{l_{i_1}, \dots, l_{i_{2L}}}$. Symbol terms encode information about the derivatives and discontinuities of the polylogarithms, and do not correspond directly to individual Feynman diagrams. More details about the map from generalized polylogarithms to symbols are given in [Appendix A](#).

The alphabet of $\mathcal{F}_{3\text{gFF}}$ is one of the simplest among all amplitudes and contains only

Transforming the Bootstrap

6

six letters, i.e., $m = 6$:

$$\mathcal{L}_{3\text{gFF}} = \{a, b, c, d, e, f\}, \quad (2)$$

cf. [Appendix A](#) These letters are Lorentz-invariant functions of the gluons' four-momenta. Via their definition they transform under a dihedral symmetry with two generators:

$$\text{cycle: } \{a, b, c, d, e, f\} \rightarrow \{b, c, a, e, f, d\}, \text{ and flip: } \{a, b, c, d, e, f\} \rightarrow \{b, a, c, e, d, f\}. \quad (3)$$

The L -loop form factor $\mathcal{F}_{3\text{gFF}}^{(L)}$ is invariant under dihedral transformations for any L .

As concrete examples, the symbols for $\mathcal{F}_{3\text{gFF}}^{(L)}$ at loops $L = 1$ and $L = 2$ contain only 6 and 12 nonvanishing terms, respectively:

$$\begin{aligned} \mathcal{S}[\mathcal{F}_{3\text{gFF}}^{(1)}] &= (-2) \left[b \otimes d + c \otimes e + a \otimes f + b \otimes f + c \otimes d + a \otimes e \right], \\ \mathcal{S}[\mathcal{F}_{3\text{gFF}}^{(2)}] &= 8 \left[b \otimes d \otimes d \otimes d + c \otimes e \otimes e \otimes e + a \otimes f \otimes f \otimes f \right. \\ &\quad \left. + b \otimes f \otimes f \otimes f + c \otimes d \otimes d \otimes d + a \otimes e \otimes e \otimes e \right] \\ &\quad + 16 \left[b \otimes b \otimes b \otimes d + c \otimes c \otimes c \otimes e + a \otimes a \otimes a \otimes f \right. \\ &\quad \left. + b \otimes b \otimes b \otimes f + c \otimes c \otimes c \otimes d + a \otimes a \otimes a \otimes e \right]. \quad (4) \end{aligned}$$

Usually, we omit the tensor product “ \otimes ” and use for example “bd” as shorthand for “ $b \otimes d$ ”, calling it a *word* or a *key*. For $L = 1$, the key bd then indexes into the tensor $C^{l_1, l_2} = C^{b, d}$, mapping onto the integer coefficient -2 . As another example, the key ab never appears in $\mathcal{S}[\mathcal{F}_{3\text{gFF}}^{(1)}]$, and therefore ab maps onto a coefficient of 0. The invariance of the form factor under dihedral transformations [\(3\)](#) relates all terms with the same coefficients in eq. [\(4\)](#) to one another. Hence at $L = 1$ ($L = 2$) there are really only 1 (2) nonzero terms to predict.

One can recover the iterated integral representation from the symbol, up to constants that can be recovered in a similar fashion. For the one-loop form factor given in eq. [\(4\)](#), its iterated integral representation is

$$\mathcal{F}_{3\text{gFF}}^{(1)} = 2 \left[\text{Li}_2(1 - bc) + \text{Li}_2(1 - ca) + \text{Li}_2(1 - ab) \right], \quad (5)$$

where $\text{Li}_2(x) = -\int_0^x dt \ln(1 - t)/t$. The iterated integral representation of the two-loop form factor $\mathcal{S}[\mathcal{F}_{3\text{gFF}}^{(2)}]$ (in a different normalization) is given in eq. (4.32) of ref. [\[35\]](#).

In general, $\mathcal{S}[\mathcal{F}_{3\text{gFF}}^{(L)}]$ is a sum of 6^{2L} *elements* (i.e., the monomials in eq. [\(4\)](#)) containing a key which is a $2L$ -sequence of letters, and an associated integer coefficient. This large number of 6^{2L} elements can be substantially reduced via a set of conditions on the symbol that restrict which letters can appear next to each other. Therefore, most of the 6^{2L} coefficients are zero, and most zeros can be accounted for by the following two simple rules:

- **adjacency rule:** any key including one of the subsequences ad, da, de (or their dihedral images) has zero coefficient [\[14\]](#).
- **prefix/suffix rule:** any key beginning with d, e or f or ending with a, b or c has zero coefficient.

Transforming the Bootstrap

We call such zero coefficients *trivial zeros*. Table 1 records the actual numbers of nonzero coefficients in the symbol at different loop orders, as well as the naive 6^{2L} counts and the counts excluding trivial zeros. The table shows that, at loop $L = 6$ and higher, about half of the terms allowed by the above two simple rules still have a zero coefficient. We refer to such terms as *nontrivial zeros*.

Loop	1	2	3	4	5	6	7	8
Total (6^{2L})	36	1,296	46,656	$1.7 \cdot 10^6$	$6.0 \cdot 10^7$	$2.2 \cdot 10^9$	$7.8 \cdot 10^{10}$	$2.8 \cdot 10^{12}$
W/O trivial zeros	6	102	1,830	32,838	589,254	$1.1 \cdot 10^7$	$1.9 \cdot 10^8$	$3.4 \cdot 10^9$
Total nonzero	6	12	636	11,208	263,880	$4.9 \cdot 10^6$	$9.3 \cdot 10^7$	$1.7 \cdot 10^9$

Table 1. Elements in the symbol $\mathcal{S}[\mathcal{F}_{3\text{gFF}}^{(L)}]$ for loops $L = 1$ to 8.

Figure 2 shows the distribution of magnitudes of nonzero integer coefficients of $\mathcal{S}[\mathcal{F}_{3\text{gFF}}]$ at loops $L = 4, 5, 6$ on a log scale. Note that the magnitudes of the coefficients grow quickly from lower to higher loops, a domain shift which may pose a challenge when attempting to train a model to generalize across loop orders.

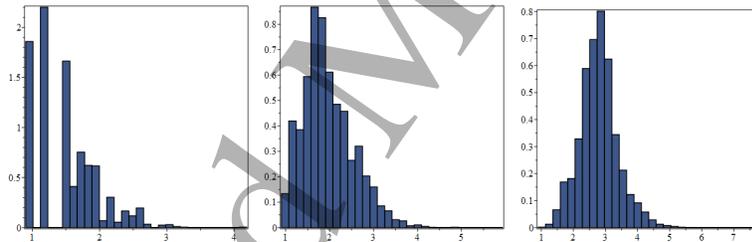


Figure 2. Histograms of the symbol coefficients for the three-gluon form factor at 4, 5, and 6 loops. The horizontal axis is the base 10 logarithm of the magnitude of the coefficient. The vertical axis is the (arbitrarily normalized) frequency with which coefficient magnitudes occur in the form factor.

2.2. Linear Relationships among Symbol Elements

In addition to the trivial zeros defined above, there exist many linear correlations in the symbol which highly constrain the values of many coefficients. We study three types of linear relations in this work—the integrability relations, the multiple-final-entry relations, and the triple-adjacency relation.

Many of these constraints are inspired by empirical observations and have deep physical roots yet to be understood. However, some are based on rather elementary mathematical considerations. For example, one important constraint is *functional integrability*: a random multi-variate symbol is not the symbol of any function, because mixed partial derivatives must commute. This requirement correlates large sets of

Transforming the Bootstrap

coefficients with specific adjacent letter pairs. One such integrability relation reads

$$F^{a,b} + F^{a,c} - F^{b,a} - F^{c,a} = 0, \quad (6)$$

which correlates the coefficients of four terms in the symbol at a time. Here $F^{a,b}$ is the abbreviation of $C^{l_1, \dots, l_{i-1}, a, b, l_{i+2}, \dots, l_{2L}}$, the coefficient corresponding to the key "...ab..." where the letter pair **ab** can appear in any pair of adjacent positions (or *slots*) in the key. The remaining letters (indicated by "...") may take any values but must be the same for each of the terms in the relation. In other words, eq. (6) can be written equivalently as

$$\begin{aligned} & C^{l_1, \dots, l_{i-1}, a, b, l_{i+2}, \dots, l_{2L}} + C^{l_1, \dots, l_{i-1}, a, c, l_{i+2}, \dots, l_{2L}} \\ & - C^{l_1, \dots, l_{i-1}, b, a, l_{i+2}, \dots, l_{2L}} - C^{l_1, \dots, l_{i-1}, c, a, l_{i+2}, \dots, l_{2L}} = 0, \end{aligned} \quad (7)$$

for any choice of $l_1, \dots, l_{i-1}, l_{i+2}, \dots, l_{2L} \in \mathcal{L}_{3\text{gFF}}$, and it remains valid for any position i in the $2L$ -length keys.

A concrete instance of relation (7) at five loops is

$$C^{\underline{c}, \underline{a}, \underline{b}, \underline{c}, \underline{a}, \underline{b}, \underline{d}, \underline{c}, \underline{c}, \underline{d}} + C^{\underline{c}, \underline{a}, \underline{c}, \underline{c}, \underline{a}, \underline{b}, \underline{d}, \underline{c}, \underline{c}, \underline{d}} - C^{\underline{c}, \underline{b}, \underline{a}, \underline{c}, \underline{a}, \underline{b}, \underline{d}, \underline{c}, \underline{c}, \underline{d}} - C^{\underline{c}, \underline{c}, \underline{c}, \underline{a}, \underline{c}, \underline{a}, \underline{b}, \underline{d}, \underline{c}, \underline{c}, \underline{d}} = 0, \quad (8)$$

where the relevant adjacent letter pairs are underlined and the four integer coefficients are 72, -88, -72, 56, respectively. We thus have $72 + (-88) - (-72) - 56 = 0$, satisfying eq. (8). Another integrability relation correlates 14 coefficients at a time, the longest linear relation currently known.

While the integrability and triple-adjacency relations occur in all adjacent pairs of slots in the key, the final-entry conditions relate only sets of keys that have the same beginnings, but different suffixes.

In general, these linear relationships can serve as an excellent probe of the learning dynamics of the models, informing us about which properties of the symbol are learned at different training stages. We explore a subset of them in Section 6, and give a more comprehensive discussion in Appendix B.

2.3. Compact Symbol Representations

With increasing loop order, the number of elements in the symbol becomes very large, as can be seen in Table 1: at 8 loops, the symbol has around 1.7 billion terms. A compact symbol representation is therefore necessary at high loops. We derive one such compact representation by noticing that the number of independent suffixes for keys of nonzero coefficients is very limited.

Applying the multiple-final-entry relations (see Appendix B) and dihedral symmetry, one notices that all terms in the symbol can be related to terms ending in the following 8 sequences of four letters: dddd, bbbd, bdbd, bbdd, dbdd, fbdd, dbbd, and cddd. We thus create a new *quad representation* where 8 new tokens are added to represent these 8 suffixes. All keys at a given loop order are then represented by their first $2L - 4$ letters plus one of the eight quad suffix letters.

Transforming the Bootstrap

9

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

Represented in the quad format, there are only 391,570 keys for the loop $L = 6$ symbol, in contrast to the 5 million keys in the original uncompressed format. Furthermore, compressing the data using the quad representation naturally removes the dihedral symmetry and all multiple-final-entry relations that involve only the last four entries. It is therefore interesting to see how the model will perform when it is presented with the more “efficient” quad format, versus having to learn these relations from the full symbol.

An even more compact *octuple representation* can be achieved by considering the last 8 letters in each element. There are 93 possible final-entry octuples after factoring out dihedral symmetry. This representation gives 16,971 keys in the $L = 6$ symbol, 312,463 in the $L = 7$ symbol, and 5.6 million in the $L = 8$ symbol. In the current study, we do not use the octuple representation, since the highest loop order under consideration is $L = 7$. However, the octuple representation may become necessary in the future to push past $L = 8$, the highest loop order currently known.

3. Implementation Details

In this section, we briefly discuss the default architecture and tokenization scheme for the later experiments.

Due to the discrete nature of our problem, all tasks are framed as sequence-to-sequence translation problems: coefficients and keys are both encoded as sequences of tokens, and the model is trained to minimize the cross-entropy of the probability distribution for the predicted coefficient sequence with the ground-truth solution. At loop L , keys are encoded as sequences of $2L$ letter tokens, e.g., ‘a, a, b, d, d, c, e, e’. While several recent works have explored different ways to tokenize integers [38, 39], we simply encode coefficients as sequences of numerical tokens in base 1000; e.g., 12334 as ‘+, 12, 334’, with the sign first [29, 40]. To preserve syntax, zero coefficients are arbitrarily assigned a sign token of ‘+’.

In most experiments, we use encoder-decoder Transformers, which contain a bidirectional Transformer encoder and an autoregressive Transformer decoder linked by a cross-attention mechanism [18]. Both encoder and decoder have the same number of layers (up to 8), the same number of attention heads (8 or 16), and the same dimension ($d = 256$, $d = 512$ or $d = 1024$). For all models, the tokenizer dimension and Transformer dimension d are the same (i.e., 256, 512, or 1024). Henceforth, we describe a model with N layers in the encoder and N layers in the decoder as an *N -layer Transformer*.

Overall, our models have between 4.5 and 245 million trainable parameters, and the best performance on many of our experiments is obtained with models with fewer than 35 million parameters. In contrast, many popular large language models [41, 42] have tens of billions of parameters. Following similar observations on AI-for-mathematics applications [23], we are able to achieve very good results with these small Transformers trained on domain-specific data. In all experiments except where noted, we use the smallest model that can perform the task with large ($> 98\%$) accuracy, where accuracy

Transforming the Bootstrap

10

is defined per element rather than per individual token.

In the default model, we use a learnable positional encoding in both the encoder and decoder [18]; alternative schemes are explored in Appendix C. The optimizer is Adam [43], with a learning rate of 10^{-4} and a flat learning rate schedule. We do not observe significant performance improvements from changes to the learning rate schedule or addition of learning-rate warmup. We believe this is due to the small size of the Transformers we use: in our experience, warmup becomes more useful for larger models.

All models are implemented in PyTorch [44] and trained on a single NVIDIA V100 GPU with 32 GB of memory, or on larger architectures (A100).

Throughout the paper, we define an epoch as a pass over 300,000 key-coefficient pairs, as opposed to the more common definition of one full pass over all training data. This makes the notion of epoch size more comparable between different experiments that employ different amounts of data at different loops. At the end of each epoch, the model is evaluated on a held-out test set. The train-test split is performed randomly. In all experiments, metrics such as accuracy are quoted on the held-out test set, not the full symbol. For example, for a training set of 100,000 elements and a test set of 1,000 elements, an accuracy of 90% refers to predicting 900 of those 1,000 elements correctly.

Accuracies are measured on a test set of N_{test} randomly sampled elements. The accuracy for each element is a binary variable (correct = 1, wrong = 0), so we may assume that the errors are binomial. Applying the central limit theorem to the average, we have a 95% confidence interval of $1.96 \times \sqrt{a(1-a)/N_{test}} \approx 2\sqrt{(1-a)/N_{test}}$ due to statistical error, for an accuracy $a \approx 1$. For $a = 99\%$ and a test set size of $N_{test} = 10,000$, this interval is 0.2%. There is also uncertainty in the accuracy due to the model's initialization.

Because the goal of these experiments is not to optimize final accuracy, we do not perform any fine-grained hyperparameter scans. For this reason, we do not typically have a separate validation set in addition to the test set. However, in order to determine whether the specific choice of training and test set might play a role, in Appendix C.2 we repeat one of our experiments for a few different train/test splits using the same set of hyperparameters chosen for the main work. We find qualitatively similar performance.

4. Predicting Symbol Coefficients from Keys

In the first experiment, we train Transformers to predict coefficients from their keys. The models are trained on a fraction of the symbol at a given loop, $L = 5$ or 6 , and are tasked to predict the remaining terms. Because most coefficients are zero, we split the problem into two separate tasks: predicting whether the coefficient corresponding to a given key is zero, and predicting a nonzero coefficient from the associated key.

We first train 1-layer Transformers with dimension $d = 256$ and 8 attention heads (i.e., 4.5 million parameters) to predict whether coefficients are zero or nonzero. We construct a dataset consisting of all nonzero-coefficient elements in the symbol plus an equal number of zero-coefficient elements. Here, zero-coefficient elements are selected

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

Transforming the Bootstrap

11

randomly from the pool of all possible zeros, and the majority of zeros are thus trivial (as defined in Section 2.1). We explore an alternative prescription to handle the nontrivial zeros in Section 6

At $L = 5$, the model correctly classifies 99.96% of elements in a test set of 10,000 examples, after only one epoch (corresponding to observing only 57% of the symbol). At $L = 6$, the model correctly classifies 99.91% of the test set after one epoch, i.e., after observing only 3% of the symbol, and 99.97% after two epochs (i.e., after observing 6% of the symbol). Distinguishing nonzero coefficients from these mostly trivial zeros thus appears to be a very simple problem for Transformers.

Predicting nonzero coefficients from their keys proves to be more difficult, necessitating larger models. We train 2-layer Transformers with dimension $d = 512$ and 8 attention heads, for loops $L = 5$ and $L = 6$. At $L = 5$, models are trained on 163,880 nonzero-coefficient elements (i.e., 62% of the symbol), and tested on 100,000 elements. At $L = 6$, models are trained on 4,816,466 nonzero-coefficient elements (i.e., 98% of the symbol), and are again tested on 100,000 elements.

At $L = 5$, the best model (out of four initializations, or *seeds*) correctly predicts 43% of the coefficients from the test set after only one epoch. Accuracy is 95% after 7 epochs, 99% after 16 epochs and 99.5% after 47 epochs. At $L = 6$, the best model (again out of four initializations) correctly predicts 95% of the coefficients in the test set after 66 epochs, 98% after 88 epochs and 99.3% after 199 epochs. We show accuracy as a function of epoch in Figure 3.

At both $L = 5$ and 6, learning proceeds in two qualitative phases: first the magnitudes of the coefficients are learned, then the signs. At $L = 5$, after two epochs, 97% of magnitudes are correctly predicted, but the signs are predicted at near chance level (50%). After 5 epochs, 99% of magnitudes are predicted correctly, but only 78% of signs are predicted correctly. By epoch 10, 98% of signs are predicted correctly. For $L = 6$, training follows the same pattern, but proceeds more slowly: the model learns the magnitudes of coefficients during the first 20 epochs. Accuracy then saturates around 50%, while the model predicts the magnitudes of coefficients with more than 95% accuracy, but predicts their signs at near chance level. Finally, from epoch 40 to 70, the model learns to correctly predict the signs of the coefficients.

Until the sign is learned, the model strongly prefers to predict one sign over the other in each epoch: predictions may for example flip from 98% ‘-’ signs in one epoch to 90% ‘+’ signs in the next. This preference gradually diminishes with epoch, decreasing to within a few percent of the true proportion of positive and negative signs (which is very close to 50% positive and 50% negative) roughly at the midpoint of the second step (epoch 7 for $L = 5$; epoch 60 for $L = 6$). We indicate this behavior for one representative run in the lower portion of Figure 3. These fluctuations are not restricted to a particular subset of magnitudes; when the true or predicted magnitude is restricted to a given value, similar gradually-decreasing fluctuations are observed.

Furthermore, when models are trained to predict only the magnitudes of nonzero coefficients at $L = 6$ (by setting all signs to ‘+’), the best model can do so at 97.7%

Transforming the Bootstrap

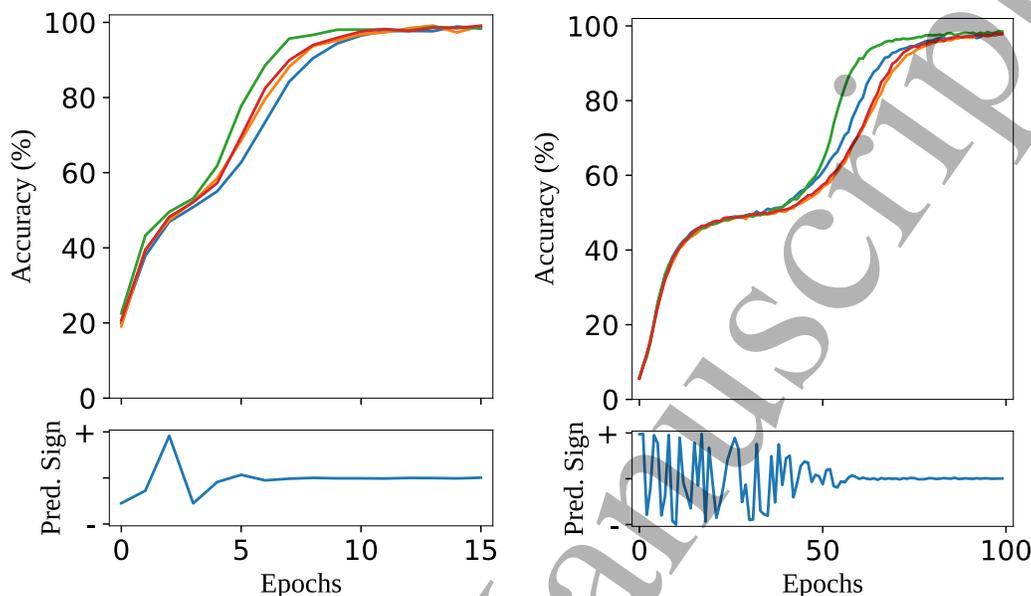


Figure 3. Accuracy vs. epoch on the nonzero coefficient-from-key task at loop $L = 5$ (left) and $L = 6$ (right), for four model initializations shown in different colors. The bottom plots show the balance of predicted signs vs. epoch, with + (−) indicating 100% (0%) positive signs. Initially the model fluctuates between strongly favoring one sign or the other before more accurately predicting the mix of signs for individual terms.

accuracy after 50 epochs. However, when models are trained to predict only the signs of the nonzero coefficients (by setting all magnitudes to ‘1’), they exhibit random guessing behavior even after 100 epochs. These results suggest that learning the magnitude of the coefficient may be a prerequisite for learning the sign.

In summary, our results indicate that Transformers trained on a small fraction of the symbol can predict coefficients from their keys with very high accuracy.

5. Quad Representation of Symbols

At higher loops, a more compact representation of the form factor symbol is necessary for efficient training. For example, the loop $L = 7$ symbol has 93 million nonzero-coefficient elements, almost 19 times as many as the $L = 6$ symbol. How long might it take to train on such a large symbol? At loop $L = 5$, performing the nonzero coefficient-from-key prediction task to $> 99\%$ accuracy takes about 22 passes through the training set and 0.7 hours. At loop $L = 6$, the same task takes 11 passes through the training set and over 54 hours. The time to reach $> 99\%$ accuracy scales at least linearly with the number of elements. If we assume this scaling continues to $L = 7$, it will take at least $54 \times 19 \approx 1000$ hours, or 43 days.

Compressing the data using the quad representation can significantly improve the

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60*Transforming the Bootstrap*

13

model training speed by reducing the number of nonzero-coefficient elements to a more manageable 7.3 million. Under the linear scaling paradigm, we would predict a more manageable 80 hours, or under 4 days.

However, learning coefficients from keys in the quad representation is a harder problem than in the full representation. The quad representation eliminates many of the obvious symmetries in the symbol: it removes both the dihedral symmetry and relations involving up to four final entries. Thus, many potential sources of correlation between the training and test sets are no longer present, and the model is forced to learn more subtle correlations between coefficients and keys in order to correctly predict the coefficients in the test set. Larger models are therefore required for this task.

Here we train models to predict nonzero coefficients from keys in the quad representation at $L = 6$ and $L = 7$. For $L = 6$, at which there are 391,570 quad keys, we use 4-layer Transformers with dimension $d = 512$ and 8 attention heads; if we train smaller 2-layer Transformers with the same dimensions, we are unable to reach above 50% accuracy even after 100 epochs of training. We train on 381,570 key-coefficient pairs and test on the 10,000 held-out elements.

The training curves for the quad representation exhibit a two-step shape very similar to those for the uncompressed representation: during the first 10 epochs, only the magnitudes of coefficients are learned and their signs are predicted at chance level. The model achieves an accuracy of 95% in 43 epochs, which equates to 34 passes over the compressed training set since an epoch is fixed at 300,000 examples. The accuracy peaks at 99% after 192 epochs, equivalent to 152 passes over the compressed training set. However, the full representation achieves 95% accuracy in 64 epochs, which equates to slightly less than 4 passes through the uncompressed training set. Thus, although the larger models trained on the quad representation can reach performance benchmarks in fewer epochs than the smaller models trained on the full representations, they in fact take more passes through the training set in order to converge, confirming our intuition that training on the quad representation is a more challenging task.

For $L = 7$, even larger models are required. 4-layer Transformers with dimension $d = 1024$ and 16 attention heads are trained on the entire quad-compressed $L = 7$ symbol (i.e., 7.3 million elements minus the held-out test set of 10,000 elements), achieving 99.1% accuracy on the test set. Here, training is considerably slower: the signs only begin to be learned after 100 epochs (*vs.* 10 epochs for the $L = 6$ symbol). The model reaches 98.5% accuracy in about 400 epochs, which equates to 120 million examples, or 16 passes over the training set. The larger training set partially accounts for this difference in learning speed.

In order to explore the relationship between model capacity and training set size, we present in Table 2 the final overall test-set accuracy for the $L = 7$ symbol in the quad representation for different model hyperparameters and training set sizes. We indicate models that fail to achieve at least 90% accuracy in gray. Accuracy decreases as training set size decreases, but remains above 94% for all but the smallest model, as long as the models are trained on 3 million examples or more. Such a training set

Transforming the Bootstrap

translates to only about 41% of the symbol elements, which suggests that the models still possess significant predictive power even when data is given in the compressed quad representation.

The Transformer’s ability to learn the sign appears to be governed largely by model capacity and training dataset size. Thus, in regions where the model sees enough data and is large enough to consistently learn the sign, accuracy is consistently $> 90\%$, while in regions where the model is too small, the sign is not learned, and accuracy stays at or below 50%. When we are near a model capacity threshold (for this task, 4 layers and $d = 512$) the model sometimes learns the sign and sometimes does not.

Arch. \ Train. size	7.3M	7M	6M	5M	4M	3M	2M	1M
8 layers, $d = 1024$	98.8%	98.7%	98.2%	97.5%	96.7%	94.8%	90.8%	78.2%
8 layers, $d = 512$	96.2%	97.4%	98.4%	96.6%	95.3%	93.8%	88.5%	36.7%
6 layers, $d = 1024$	98.6%	98.9%	98.0%	97.9%	96.7%	94.8%	90.3%	58.5%
6 layers, $d = 512$	95.2%	96.6%	96.9%	95.8%	94.4%	94.5%	87.9%	34.8%
4 layers, $d = 1024$	99.1%	98.9%	98.3%	97.9%	96.6%	94.9%	89.9%	39.1%
4 layers, $d = 512$	48.5%	96.0%	94.1%	48.3%	94.6%	81.7%	55.3%	33.9%

Table 2. Maximum test-set accuracy after 250 epochs at loop 7 in the quad representation, for various training set sizes as well as numbers of layers and dimensions. The best of two models is shown. All models have 16 heads. (The smallest model occasionally does not emerge from the first plateau, with its accuracy then staying below 50%.)

6. Model Characterization via Relationship Accuracy

The results of the previous experiments strongly suggest that the model leverages certain correlations that are present in the data, such as dihedral symmetry and the final-entry relations described in Section 2.2, in order to more easily extrapolate from the training set into the test set. Additionally, the unusual two-phase accuracy curves that occur in both the full and quad representations warrant further investigation. In this section, we therefore explore how the linear relations behave as a function of epoch in order to better understand how the model learns.

We define an *instance* of a relation as a set of keys and their associated coefficients that obey a given relation. For example, for the relation given in eq. (6), a sample instance at 5 loops is, as in eq. (8):

$$C_{c,a,b,c,a,b,d,c,c,d} + C_{c,a,c,c,a,b,d,c,c,d} - C_{c,b,a,c,a,b,d,c,c,d} - C_{c,c,a,c,a,b,d,c,c,d} = 0 \quad (9)$$

which corresponds to the following set of keys and coefficients: {cabcabdccc: 72, caccabdccc: -88, cbcabdccc: -72, ccabdccc: 56 }.

We generate 500 instances of each homogeneous linear relation at loop L and use them to evaluate the performance of a model trained on the coefficient-from-key

Transforming the Bootstrap

15

prediction task as in Section 4. To do so, we randomly generate a set of keys that obey the given relation and then pair them with the corresponding coefficients. We discard and re-generate all such multi-term instances that do not contain at least one nonzero coefficient. In this way, we avoid instances with all zero terms that are trivially satisfied.

The relation instances generated are used only as an auxiliary test set. The training set consists of the full nonzero symbol plus an equal proportion of zeros as in Section 4 while the coefficient-from-key test set is still employed as before. We note that more than 99% of the nonzero terms in the relation instances appear in the training set. However, we stress that this does not constitute data leakage, as the relation evaluations are auxiliary dataset-level metrics that do not influence training.

The linear relations may relate nonzero terms to the nontrivial zeros, which constitute a very small fraction of all possible zeros. Therefore, when choosing the zeros to be added to the training set, we explicitly select a large proportion of nontrivial zeros. The fraction of trivial zeros in the training set is restricted to be 5% of all zeros. Prioritizing the nontrivial zeros causes accuracy to decrease slightly on the trivial zeros; however, as trivial zeros are easy to learn and to identify (as per Section 4) we can simply manually set the predicted coefficient to zero for any trivial-zero terms in a relation instance.

The model used for this experiment is a 2-layer Transformer with $d = 512$ and 8 heads, trained on 9,732,932 elements (i.e., the full $L = 6$ symbol plus an equal proportion of zeros), and tested on 100,000 randomly chosen held out elements. After 200 epochs, the model correctly predicts 98.47% of the coefficients in the test set. The learning curves again reveal the familiar two qualitative phases, though training to a given accuracy now takes twice as many epochs (due to the addition of zeros increasing the dataset size by a factor of two) and overall magnitude and sign accuracy both reach 50% accuracy within the first epoch, due to the fact that zeros are learned quickly.

A complete list of the relations we evaluate is given in Appendix B. Here, we only describe the following short relations, which form a representative subset of the different types of relations studied. In the triple and integrability relations, the specified adjacent slots can appear anywhere in the key, while the specified adjacent slots in the final-entry relations must appear at the end of the key (which we denote by \mathcal{E} instead of F).

$$\text{triple 0: } F^{a,a,b} + F^{a,b,b} + F^{a,c,b} = 0, \quad (10)$$

$$\text{integ 0: } F^{a,b} + F^{a,c} - F^{b,a} - F^{c,a} = 0, \quad (11)$$

$$\text{integ 1: } F^{c,a} + F^{c,b} - F^{a,c} - F^{b,c} = 0, \quad (12)$$

$$\text{final 16: } \mathcal{E}^{b,f} - \mathcal{E}^{b,d} = 0, \quad (13)$$

$$\text{final 17: } \mathcal{E}^{c,d,d} + \mathcal{E}^{c,e,e} = 0, \quad (14)$$

$$\text{final 18: } \mathcal{E}^{d,d,b,d} - \mathcal{E}^{d,b,d,d} = 0. \quad (15)$$

For each relation, we quote four metrics: 1) whether the coefficients predicted by the model satisfy the given relation, regardless of whether the individual coefficients themselves are correct (red); 2) whether the relation is satisfied and all coefficients in

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

Transforming the Bootstrap

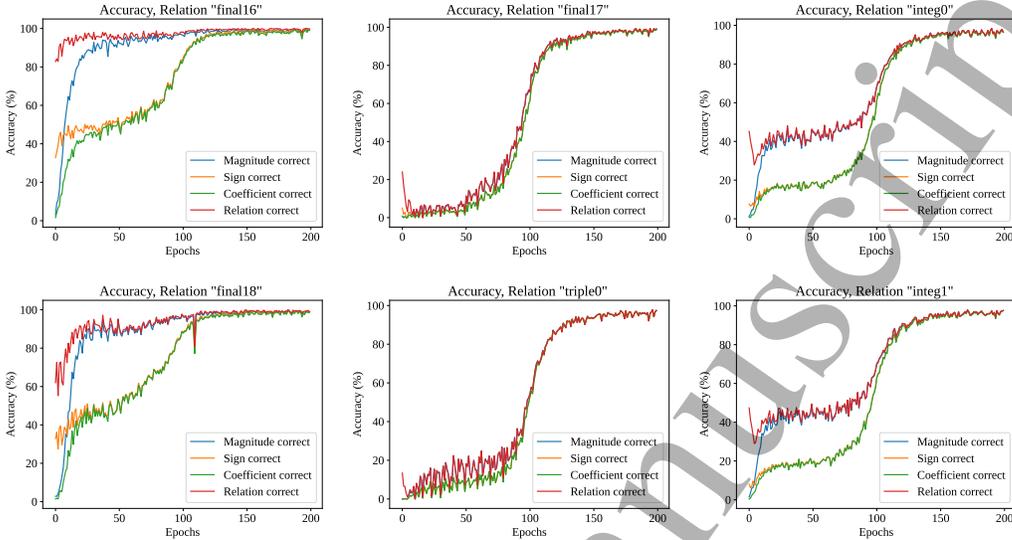


Figure 4. Relation accuracy (red), magnitude accuracy (blue), sign accuracy (yellow), and coefficient accuracy (green) for each of the named relations, grouped by behavior. Relations in Group 1 (left column) are two-term equivalence relations that are consistently satisfied after only a few epochs. Relations in Group 2 (center column) are relations that require at least two coefficients to have different signs, and are not satisfied until the second phase. Relations in Group 3 (right column) are mixed relations, for which some instances decompose into pairs of equivalent terms (as in Group 1 relations) while others do not.

the instance have the correct magnitudes, regardless of their signs (blue); 3) whether the relation is satisfied and all coefficients in the instance have the correct signs, regardless of their magnitudes (yellow); and 4) whether all coefficients in the instance are correct (green). We plot these metrics as a function of epoch in Figure 4. Similar results for additional relations are provided in Appendix B

The linear relations can be grouped by structure into three categories that also define their behavior. The six relations discussed in this section contain two examples from each category.

- Group 1 relations, such as eq. (13) (final 16) and eq. (15) (final 18), are equivalence relations that require two coefficients to have the same magnitudes and signs. For all Group 1 relations, the relation is often satisfied before all the magnitudes are predicted correctly—the model predicts that both coefficients in the instance must be the same before it is able to successfully identify what that coefficient is. While the model learns the magnitudes of both coefficients in the relation instances fairly quickly, the signs of the coefficients are only predicted correctly 50% of the time until the sign is learned; however, both coefficients are consistently predicted to have the same sign. Thus, the sign fluctuations described in Section 4 (Figure 3) largely respect the Group 1 relations.

Transforming the Bootstrap

17

- Group 2 relations, such as eq. (10) (triple 0) and eq. (14) (final 17), are those in which at least two coefficients must have opposite signs. While the magnitudes of coefficients in these relation instances are learned within a few epochs, accuracy on all relation metrics remains low, but steadily increasing, until the signs are learned. This behavior is also largely dictated by the sign fluctuations shown in Figure 3.
- Group 3 relations, such as eq. (11) (integ 0) and eq. (12) (integ 1), are multi-term relations that may be satisfied by two or more pairs of identical coefficients or by a set of related but nonidentical coefficients (e.g., 40% of the generated (integ 0) instances are expressible as pairs of identical coefficients). Under the conditions where the model has a high probability of satisfying the Group 1 relations, the model predictions for the subset of Group 3 relation instances expressible as pairs of identical coefficients will also satisfy the relation.

These properties suggest an explanation for the double plateau behavior: first, the model learns to group elements whose coefficients have the same magnitude; then it learns to correctly predict those magnitudes. The model predictions for the sign fluctuate from epoch to epoch—rather wildly at first—until the second accuracy step is reached. However, these fluctuations consistently respect the Group 1 equivalence relations between elements. During this fluctuation phase, the model also gradually learns the Group 2 relations; i.e., the model learns which coefficients with a given magnitude have the same sign and which do not. These fluctuations get smaller until the sign is eventually learned.

Both the cycle and flip symmetries can be expressed as Group 1 relations relating pairs of terms with identical coefficients. We plot the relation evaluation metrics for these relations in Appendix B and find that they behave similarly to other Group 1 relations.

We observe a further intriguing manifestation of the dihedral symmetry in the geometry of the embedding layer representation. Specifically, we extract the learned d -dimensional embeddings of the input letter tokens from the embedding layer of the Transformer and calculate the angles between them. These embedding vectors obey both the cycle and flip symmetries. We perform such an extraction experiment with a 2-layer Transformer with $d = 512$, first at $L = 5$ and then at $L = 6$.

At $L = 5$, the triangles $\triangle abc$ and $\triangle def$ are approximately equilateral: all angles are within 1.5° of 60.0° for $\triangle abc$ and within 2.7° of 60.0° for $\triangle def$. This result indicates that the embedding vectors obey the cycle symmetry. Similarly, all angles are within 3.7° of 60.0° for triangle $\triangle abf$, within 1.5° of 60.0° for $\triangle bcd$, and within 3.3° of 60.0° for $\triangle ace$; the fact that these triangles are approximately similar indicates that the embedding vectors obey the flip symmetry.

At $L = 6$, we observe the same phenomenon even more strongly: $\triangle abc$ and $\triangle def$ are approximately equilateral: all angles are within 1.0° of 60.0° for $\triangle abc$ and within 0.6° of 60.0° for $\triangle def$, indicating cycle symmetry. Likewise, all angles are within 0.5° of 60.0° for triangle $\triangle abf$, within 0.7° of 60.0° for $\triangle bcd$, and within 0.8° of 60.0° for $\triangle ace$,

Transforming the Bootstrap

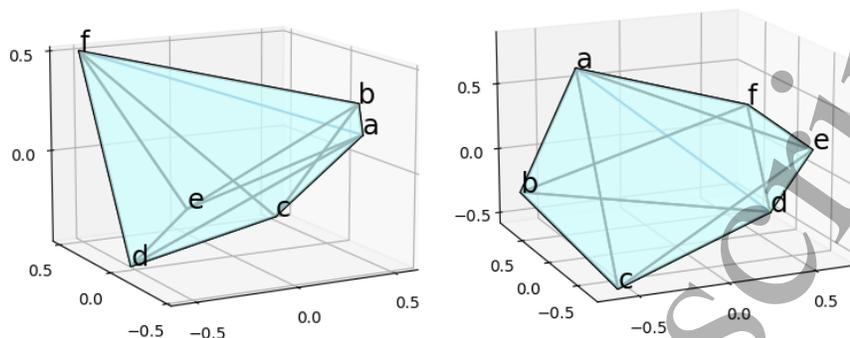


Figure 5. (Left) The leading three PCA components of token embeddings for a 2-layer Transformer with $d = 512$ trained for 50 epochs on $L = 5$ data, with zeros included. The leading three PCA components explain 63.56% of variance, and dihedral symmetry is not visible. (Right) The leading three PCA components of token embeddings for a 2-layer Transformer with $d = 512$ trained for 200 epochs on $L = 6$ data, with zeros included. The leading three PCA components explain 81.76% of variance. The octahedron exhibits dihedral symmetry.

indicating flip symmetry.

We perform standard linear principal component analysis (PCA) and plot the embeddings of these letter tokens in the space of the three leading PCA components in Figure 5. Projecting to the leading three components distorts the angles somewhat: at $L = 5$, the dihedral symmetry is no longer apparent, while at $L = 6$, the cycle and flip symmetries are visually apparent but the octahedron is no longer regular.

7. Mixed-loop Training

To successfully extend the bootstrap program to unseen loops, we must build models that can generalize from lower loops, for which we have the complete symbol, to higher loops, where only a small number of symbol terms may be available.

However, in many AI-for-mathematics applications, Transformers trained exclusively at one input length using absolute position encoding fail to generalize to different input lengths [39]. Here we face a similar challenge. When using loop $L = 6$ data to evaluate a model trained exclusively at $L = 5$ (and vice versa) for the task of predicting nonzero coefficients, our baseline models can only attain an accuracy of at most 3% for a variety of model sizes and depths. Many predicted coefficients at unseen loops are nonsensical, such as the string ‘+++’. Given that our ultimate goal is to predict coefficients of keys at unseen higher loops, this failure of length generalization presents a major limitation that we must overcome.

In many ways, however, this problem goes even beyond simple length generalization. At each subsequent loop, the number of possible values of keys and coefficients both increase appreciably (see Figure 2), and it is not clear whether or how the functional

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

Transforming the Bootstrap

form that relates them changes as well. Therefore, while alternative architecture designs, positional encoding schemes, and numerical encoding schemes (which we explore in [Appendix C](#)) may be helpful for this task, a much more comprehensive strategy will likely be needed.

As a first attempt to address this issue, we train Transformers to predict coefficients using an even proportion of $L = 5$ and $L = 6$ data. The full nonzero $L = 5$ symbol (263,880 elements) is first combined with an equal proportion of zero-coefficient elements at $L = 5$. Another 263,880 nonzero elements are then drawn from the $L = 6$ symbol (representing 5% of the $L = 6$ symbol) and augmented with a roughly equal amount of zero-coefficient elements at $L = 6$. The zero-sampling is done naively, as in [Section 4](#) rather than in the nontrivial-zero-biased manner of [Section 6](#). Training and test sets are constructed for $L = 5$ and $L = 6$ separately, and they are then merged to create *mixed-loop* training and test sets. Each loop-specific training set for both $L = 5$ and $L = 6$ contains 517,760 examples, while each loop-specific test set contains 10,000 examples; every set is an equal mix of zero- and nonzero-elements. Additionally, we create a larger $L = 6$ training set that is the same size as the mixed-loop set (1,035,520 examples,

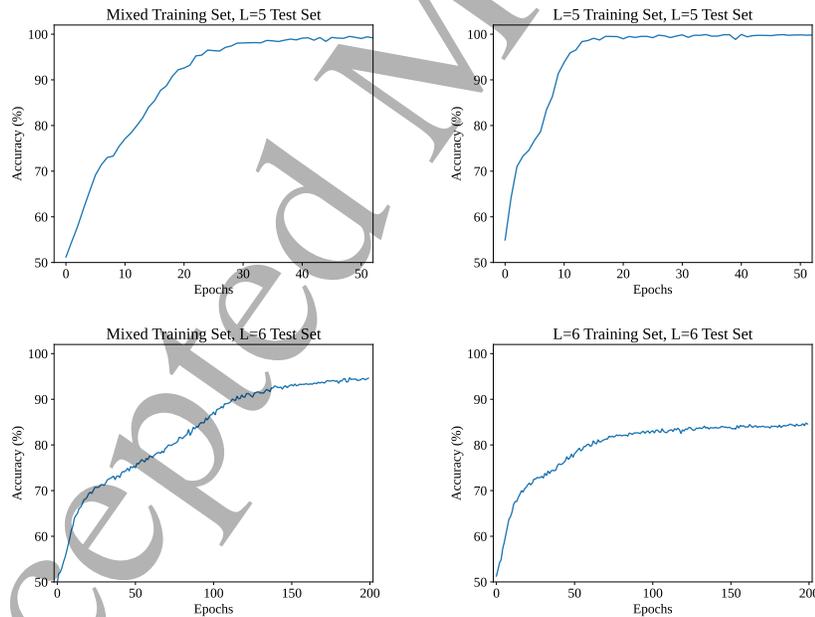


Figure 6. Accuracy vs. epoch for models trained on the mixed and per-loop training sets, evaluated on the per-loop test sets. In all cases, accuracy starts at 50% because zeros are learned in the first epoch. Models trained on the mixed training set take almost exactly half the number of epochs to reach performance benchmarks for $L = 5$ as models trained on the dedicated $L = 5$ training set, corresponding to approximately equal performance of these two models. However, models trained on a small subset of the $L = 6$ symbol mixed with the full $L = 5$ symbol are able to generalize to the $L = 6$ test set much better than models trained on the $L = 6$ symbol alone.

Transforming the Bootstrap

20

roughly 10% of the $L = 6$ symbol), in order to evaluate whether the effects of mixed-loop training can be explained by the difference in training set size.

We train a model with 2 layers, $d = 512$, and 8 attention heads in both the encoder and decoder for 200 epochs, evaluating on both the mixed-loop test set and the individual $L = 5$ and $L = 6$ test sets. The model again exhibits two-phase learning behavior in all cases, as shown in Figure 6 (which displays results for all but the larger $L = 6$ training set). We measure the performance by reporting three epochs,

- (i) the epoch at which the test-set magnitude accuracy first exceeds 90%,
- (ii) the midpoint epoch of the plateau step, which is when the model’s overall accuracy first reaches the average between its final-state accuracy and 75%,
- (iii) the epoch at which the overall test-set accuracy first exceeds 90%,

as well as the best overall test-set accuracy after 200 epochs.

In Table 3 we compare the model performance when trained on each loop-specific training set to its performance when trained on the mixed training set. The dedicated $L = 5$ model reaches all performance benchmarks for each loop in almost exactly half the number of epochs as the mixed-loop model. Accounting for the fact that the mixed-loop training set is exactly twice the size of the individual $L = 5$ and $L = 6$ training sets, this result suggests that the mixed-loop model solves the prediction problem for $L = 5$ in almost exactly the same amount of time as it would without the addition of $L = 6$ data. On the other hand, the mixed-loop model learns the $L = 6$ magnitudes in approximately the same number of epochs as the dedicated $L = 6$ model, corresponding to reaching this benchmark (i.e., First Mag. Acc. $> 90\%$) in half as many iterations through the $L = 6$ training set. The mixed-loop model also exceeds $\sim 90\%$ overall accuracy, whereas the dedicated $L = 6$ model does not. However, much of this effect may be explained by the dataset size: when a model is trained on the larger $L = 6$ dataset that is the same size as the mixed set, the results appear highly similar to those we see when using the mixed training set.

Train	Eval	First Mag. Acc. > 90% [Epoch]	Midpoint of Step [Epoch]	First Total Acc. > 90% [Epoch]	Best Acc. Epoch 200
Mix	$L = 5$	6	16	18	99.82%
$L = 5$	$L = 5$	3	9	10	99.9%
Mix	$L = 6$	25	88	112	94.6%
$L = 6$	$L = 6$	20	51	N/A	84.5%
$L = 6$, Large	$L = 6$	17	92	100	97.67%

Table 3. Model learning dynamics for mixed training: the epoch at which the test-set magnitude accuracy first exceeds 90%, the midpoint of the plateau step, the epoch at which the overall test-set accuracy first exceeds 90%, and the best overall test-set accuracy after 200 epochs. All training hyperparameters (architecture, initialization seed, etc.) are kept the same; only the training and test sets are changed between runs.

Transforming the Bootstrap

21

These results are nonetheless rather encouraging: while adding a subset of $L = 6$ data to $L = 5$ does not lead to improvement on $L = 5$ tasks, adding the full $L = 5$ symbol to a subset of the $L = 6$ symbol does appear to improve performance on $L = 6$ tasks. This suggests that features learned at lower loops can be employed to enhance the predictive power of a small number of symbol elements at higher loops (for which we may only be able to determine a handful of coefficients *a priori*).

8. Steps Toward Predicting the Next Loop

Although mixed-loop training may allow us to better generalize from only a small number of symbol elements at unseen loops, we still wish to find ways to predict these elements at loops for which we do not have coefficient information.

As a first attempt at solving this task, we now consider a different problem setup. Instead of predicting unknown symbol elements at a given loop order, we would like to obtain the loop L symbol from the coefficients at loop $(L - 1)$. In other words, for any element at loop L , we want to recover its coefficient from the coefficients of a list of *parent* elements at one loop lower that are related by having similar strings of letters in their keys.

The keys of elements at loop L are sequences of $2L$ letters, whereas the keys at loop $L - 1$ are only $(2L - 2)$ letters long. We therefore define the *strike-two parents* of a given key at loop L as the keys from loop $(L - 1)$ which are created by simply striking out two letters from the key. For instance, the six strike-two parents of the loop $L = 2$ key `aacf` are

$$\overset{\cdot}{a}\overset{\cdot}{a}cf=cf, \overset{\cdot}{a}\overset{\cdot}{a}cf=af, \overset{\cdot}{a}\overset{\cdot}{a}cf=ac, \overset{\cdot}{a}\overset{\cdot}{a}cf=af, \overset{\cdot}{a}\overset{\cdot}{a}cf=ac, \overset{\cdot}{a}\overset{\cdot}{a}cf=aa, \quad (16)$$

which are keys at $L = 1$. In general, there are $\binom{2L}{2} = L(2L - 1)$ strike-two parents of any key at loop L .

The dataset for this experiment is constructed by first selecting certain elements at loop L . For each key, we then construct the list of strike-two parents at loop $L - 1$ and list their coefficients in the *strikeout order*, i.e., the coefficient corresponding to striking the first two letters from the key is ordered first in the coefficient list, the coefficient corresponding to striking the first and third letters from the key is ordered second in the list, etc. For the above example at $L = 2$, the coefficient to be predicted is $C^{aacf} = 0$, and the ordered list of its strike-two parent coefficients is $[0, -2, 0, -2, 0, 0]$. Many of the parent coefficients are zero due to the nature of the strike-out. For example, pairs of letters that are not allowed to be adjacent can become adjacent in the parent keys after a letter in between is removed.

In these experiments, 4-layer Transformers with $d = 512$ and 8 heads are trained to predict nonzero coefficients at $L = 6$ from the coefficients of their $L = 5$ parents. Model inputs are sequences of $\binom{12}{2} = 66$ parent coefficients at $L = 5$ given in strikeout order, while the targets are single coefficients at $L = 6$.

Transforming the Bootstrap

22

From the $L = 5$ and $L = 6$ symbols, we can create 4.9 million examples ($L = 5$ parents and $L = 6$ coefficients), but this dataset includes many duplicates. In fact, each example is duplicated 6.4 times on average, primarily due to the dihedral symmetry. To avoid contamination between the training and test set, we restrict the dataset to the 767,500 unique examples, split into 757,500 training and 10,000 test examples.

We report our results in the first row of Table 4. After 500 epochs, the model predicts 98.1% of the test examples. Learning is fast: 90% accuracy is achieved after 20 epochs, and 95% after 80 epochs. In these experiments, the magnitudes and signs of coefficients are learned simultaneously. In other words, no two-phase learning dynamics are observed, in contrast to previous experiments.

Our results suggest that there may exist learnable formulas for computing coefficients at $L = 6$ from their strike-two parents at $L = 5$. We are not, so far, capable of explicitly recovering these formulas, but additional ablation experiments may shed light on some of their features. We perform some such experiments in the remainder of this section, and give further results in Appendix D.

First, we investigate whether we can predict the $L = 6$ coefficient from a smaller set of $L = 5$ parents. One way is to only strike letters that are no more than k positions away from each other in the $L = 6$ keys. For example, with $k = 1$ we only strike adjacent letters; for $k = 2$ we only strike letters that are either adjacent or separated by one additional letter, etc. The number of parents remaining in this reduced set is $2kL - \frac{k(k+1)}{2}$. At $L = 6$, the 66 parents available at the maximum $k = 12$ are reduced to only 11 parents for $k = 1$. We construct the reduced dataset for $L = 6$ with $k = 1, 2, 3, 5$, remove duplicates, and present the parent coefficients to the model in the knockout order.

Our models predict 98.3%, 98.4%, 98.1% and 94.3% of test examples for $k = 5, 3, 2$ and 1, respectively. Predicting from 21 parents ($k = 2$), instead of the full 66, has little impact on model performance, suggesting that the majority of coefficients at higher loops may be learnable using only a limited set of parent coefficients at lower loops.

Next we experiment with the order of the knockout parents. Surprisingly, even if parent coefficients (i.e., model inputs) are randomly shuffled, the model can still achieve 95.2% accuracy (93.5% with the 21 parents for $k = 2$). When parents are sorted by increasing order in their numerical values, the model achieves 93.9% for $k = 5$. This result suggests that the ability of Transformers to compute coefficients at $L = 6$ from $L = 5$ is mostly unaffected by permutations of the $L = 5$ coefficients. However, the result may be an artifact of the way the knockout experiment is constructed. Because certain letter adjacency conditions lead to zero coefficients, an ordered list of parents implicitly encodes some information about the original letter structure of the key. From previous experiments, we know that information about the letter structure of the key can be used to reconstruct the coefficient. In future knockout experiments, we may wish to take this effect into account in order to better model the relationships between elements across loop orders.

Finally, we modify the values of the parent coefficients themselves in two different ways: (1) we set the signs of all parents to ‘+’ before removing duplicates, thereby

	Accuracy	Magnitude accuracy	Sign accuracy
Strike-two, all parents	98.1%	98.4%	99.6%
Strike-two, $k = 5$	98.3%	98.6%	99.7%
Strike-two, $k = 3$	98.4%	98.7%	99.7%
Strike-two, $k = 2$	98.1%	98.3%	99.5%
Strike-two, $k = 1$	94.3%	95.2%	98.5%
Randomly shuffled parents, all parents	95.2%	99.1%	96.3%
Randomly shuffled parents, $k = 2$	93.5%	98.1%	95.0%
Sorted parents, $k = 5$	93.9%	95.4%	97.9%
Parent magnitudes only, all parents	81.8%	98.4%	83.2%
Parent signs only, all parents	93.3%	93.5%	99.0%
Parent signs only, all parents, sorted	0.8%	61.0%	1.6%

Table 4. Overall, magnitude, and sign accuracy for the cross-loop strike-out experiments described in the text. Best of four seeds, trained for about 500 epochs.

retaining only the magnitude information; or (2) we provide only the signs of the parent coefficients as inputs to the model, i.e., each parent coefficient is encoded as -1 , $+0$ or $+1$. In both cases, all modified parent coefficients are still presented in the knockout order.

In case (1), models trained on only the magnitudes of the parents are able to recover the magnitudes of the target coefficients with 98.4% accuracy, about the same level as models trained on unmodified parent coefficients. However, the sign of the target coefficient proves harder to learn when the model cannot see the signs of the parents (dropping from 99.6% to 83.2% accuracy).

In case (2), the models can achieve an overall accuracy of 93.3%, and correctly predict the sign in 99% of the test cases. However, if we additionally shuffle or sort these parents in ascending order (i.e., all -1 s, all 0 s, then all $+1$ s), we find that the model is totally unable to learn. In other words, we can drastically reduce information about either the values of the knockout parents or their ordering and still recover the full coefficient, but we cannot do both simultaneously.

The fact that coefficients can be reconstructed reliably from their strike-two parents, despite the fact that symmetries such as dihedral symmetry are removed, suggests that a closed-form solution for computing coefficients from their parents may exist. It is likely that some amount of redundant information exists in the set of knockout parents, as we are able to reconstruct the coefficient in a number of scenarios when the set of parents is severely altered. We plan to investigate such cross-loop relationships further in future work.

9. Conclusions and Future Work

In this study, we have shown that a Transformer model is able to successfully predict the coefficients of elements in the symbols for scattering amplitudes in $\mathcal{N} = 4$ planar super Yang-Mills theory. Below we summarize the key findings of our work.

Our models learn in a two-phase fashion, first achieving very high accuracy on the coefficient magnitudes and then learning their signs. The distribution of predicted signs exhibits large fluctuations that gradually decrease as the sign is learned. The models cannot learn the signs of coefficients without the magnitude information, while they can learn magnitudes without sign information.

Transformers perform very well even when the data is compressed into the quad representation, where many trivial correlations between terms are removed. Enough information remains for a (larger) Transformer to successfully reconstruct coefficients from keys. This result bodes well for our ability to move to higher loops, as the space of possible coefficients and keys becomes quite large beyond $L = 6$.

To study the learning dynamics, we have evaluated the models' performance on the linear relations. The relations can be classified into three groups based on whether they require coefficients to be identical, to have opposite signs, or some combination of the two. We propose a likely explanation for the two-phase behavior by assessing which of the known linear relations are satisfied at each epoch: the model first learns to group many terms with the same coefficient magnitude; then it learns what those magnitudes are; next it learns which coefficients with a given magnitude have the same sign and which have a different sign; finally it learns the true signs. Additionally, one of the simplest symmetries of the symbol, its dihedral symmetry, can be seen geometrically in the embedding layers of the Transformer.

We have also trained models on mixed-loop data. Augmenting a small percentage of training examples at one loop with a substantial fraction of the symbol at lower loops leads to faster convergence and higher accuracy on a test set at the higher loop. This performance is particularly encouraging, as it suggests that only a relatively small number of coefficients may need to be provided at unseen loops in order for Transformers to successfully predict the rest of the symbol.

In our second set of experiments, namely the cross-loop strike-out approach, we have shown that a model can predict the coefficients at loop L using only a small subset of related coefficients at loop $(L - 1)$. Coefficient information from the lower loop can be scrambled or severely degraded (but not both at once), without hindering the ability to reconstruct the target coefficient.

Transformers are well-studied, but they are not the only architecture with the potential to learn the structure of scattering amplitudes. In [Appendix C.1](#), we train a Long-Short-Term Memory recurrent neural network (LSTM) [\[45\]](#) and a Gated Recurrent Unit (GRU) [\[46\]](#) on the $L = 5$ symbol. In both cases, we are able to attain comparable performance to a Transformer. It may be interesting to explore other types of architectures as well.

Transforming the Bootstrap

25

In the future, we plan to train on more complicated objectives, with the goal of developing a model that encodes many types of information from multiple loops. This evokes the concept of a “foundation model”: by training a large, multitask model with information about a number of relevant concepts, we hope to better characterize the complex recurrences and relations present in the symbol in order to generalize to unseen loops. This is a challenging domain-generalization task, as the distribution of possible keys and coefficients changes substantially between loop orders.

The methods in this paper can be applied straightforwardly to other SYM problems where there is multi-loop data encoded by symbols. Additional challenges arise when considering a similar approach to amplitudes in QCD. While many of the generalized polylogarithms appearing in SYM also appear in QCD, the QCD result does not have uniform weight $2L$, but generally has all weights from $2L$ down to zero (rational functions). The lower-weight terms typically have complicated prefactors that are rational functions of the kinematic variables. Also, the behavior in physical limits is not understood as well in QCD as in SYM. A better understanding of how to organize and bootstrap QCD amplitudes would clearly benefit any machine-learning approaches to determining them.

Ultimately, our goal is to build machine-learning models capable of computing amplitudes analytically *ab initio*. In addition to the potential spin-offs for collider physics, understanding even simplified scattering amplitudes to all loop orders would give a remarkable new window into quantum field theory.

Acknowledgments

The datasets used in this paper are available at <https://doi.org/10.5281/zenodo.11218272>. A prior version of this work was presented in a non-archival form at the 2023 NeurIPS Machine Learning for the Physical Sciences workshop [47]. We thank Yang-Hui He, Romuald Janik, Matt Schwartz, Jesse Thaler, Dimitris Papailiopoulos, Jordan Ellenberg, Gary Shiu, and Yiqiao Zhong for fruitful discussions. We would like to thank Phil Wang (“lucidrains”) for use of his implementation of the RoPE and xPos relative position encodings. The work was supported in part by the U.S. Department of Energy (DOE) under Award No. DE-FOA-0002705, KA/OR55/22 (AIHEP). LD and TC are additionally supported by the U.S. Department of Energy Award No. DE-AC02-76SF00515. MW was supported by the research grant 00025445 from Villum Fonden.

References

- [1] A. Huss, J. Huston, S. Jones and M. Pellen, *Les Houches 2021—physics at TeV colliders: report on the standard model precision wishlist*, *J. Phys. G* **50** (2023) 043001 [2207.02122].
- [2] G. Heinrich, *Collider Physics at the Precision Frontier*, *Phys. Rept.* **922** (2021) 1 [2009.00516].
- [3] C. Anastasiou, C. Duhr, F. Dulat, F. Herzog and B. Mistlberger, *Higgs Boson Gluon-Fusion Production in QCD at Three Loops*, *Phys. Rev. Lett.* **114** (2015) 212001 [1503.06056].
- [4] C. Anastasiou, C. Duhr, F. Dulat, E. Furlan, T. Gehrmann, F. Herzog et al., *High precision determination of the gluon fusion Higgs boson cross-section at the LHC*, *JHEP* **05** (2016) 058 [1602.00695].

Transforming the Bootstrap

26

- [5] B. Mistlberger, *Higgs boson production at hadron colliders at N^3LO in QCD*, *JHEP* **05** (2018) 028 [1802.00833].
- [6] C. Duhr, F. Dulat and B. Mistlberger, *Higgs Boson Production in Bottom-Quark Fusion to Third Order in the Strong Coupling*, *Phys. Rev. Lett.* **125** (2020) 051804 [1904.09990].
- [7] F. A. Dreyer and A. Karlberg, *Vector-Boson Fusion Higgs Production at Three Loops in QCD*, *Phys. Rev. Lett.* **117** (2016) 072001 [1606.00840].
- [8] C. Duhr, F. Dulat and B. Mistlberger, *Drell-Yan Cross Section to Third Order in the Strong Coupling Constant*, *Phys. Rev. Lett.* **125** (2020) 172001 [2001.07717].
- [9] C. Duhr, F. Dulat and B. Mistlberger, *Charged current Drell-Yan production at N^3LO* , *JHEP* **11** (2020) 143 [2007.13313].
- [10] C. Duhr and B. Mistlberger, *Lepton-pair production at hadron colliders at N^3LO in QCD*, *JHEP* **03** (2022) 116 [2111.10379].
- [11] L. J. Dixon, J. M. Drummond and J. M. Henn, *Bootstrapping the three-loop hexagon*, *JHEP* **11** (2011) 023 [1108.4461].
- [12] S. Caron-Huot, L. J. Dixon, A. McLeod and M. von Hippel, *Bootstrapping a Five-Loop Amplitude Using Steinmann Relations*, *Phys. Rev. Lett.* **117** (2016) 241601 [1609.00669].
- [13] S. Caron-Huot, L. J. Dixon, F. Dulat, M. von Hippel, A. J. McLeod and G. Papathanasiou, *Six-Gluon amplitudes in planar $N = 4$ super-Yang-Mills theory at six and seven loops*, *JHEP* **08** (2019) 016 [1903.10890].
- [14] L. J. Dixon, Ö. Gürdoğan, A. J. McLeod and M. Wilhelm, *Bootstrapping a stress-tensor form factor through eight loops*, *JHEP* **07** (2022) 153 [2204.11901].
- [15] T. Gehrmann, M. Jaquier, E. W. N. Glover and A. Koukoutsakis, *Two-Loop QCD Corrections to the Helicity Amplitudes for $H \rightarrow 3$ partons*, *JHEP* **02** (2012) 056 [1112.3554].
- [16] B. Basso, A. Sever and P. Vieira, *Spacetime and Flux Tube S-Matrices at Finite Coupling for $N=4$ Supersymmetric Yang-Mills Theory*, *Phys. Rev. Lett.* **111** (2013) 091602 [1303.1396].
- [17] A. Sever, A. G. Tumanov and M. Wilhelm, *Operator Product Expansion for Form Factors*, *Phys. Rev. Lett.* **126** (2021) 031602 [2009.11297].
- [18] A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez et al., *Attention Is All You Need*, in *31st International Conference on Neural Information Processing Systems*, 2017, [1706.03762].
- [19] D. Bahdanau, K. Cho and Y. Bengio, *Neural Machine Translation by Jointly Learning to Align and Translate*, in *3rd International Conference on Learning Representations*, 2015, [1409.0473].
- [20] J. Devlin, M. Chang, K. Lee and K. Toutanova, *BERT: pre-training of deep bidirectional transformers for language understanding*, in *17th Annual Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 2019, [1810.04805].
- [21] A. Dosovitskiy, L. Beyer, A. Kolesnikov, D. Weissenborn, X. Zhai, T. Unterthiner et al., *An image is worth 16×16 words: Transformers for image recognition at scale*, in *9th International Conference on Learning Representations*, 2021, [2010.11929].
- [22] B. Romera-Paredes, M. Barekatin, A. Novikov, M. Balog, M. P. Kumar, E. Dupont et al., *Mathematical discoveries from program search with large language models*, *Nature* **625** (2024) 468.
- [23] F. Charton, *Linear algebra with transformers*, *Trans. Mach. Learn. Res.* **2022** (2022) [2112.01898].
- [24] S. Caron-Huot, L. J. Dixon, J. M. Drummond, F. Dulat, J. Foster, O. Gürdoğan et al., *The Steinmann Cluster Bootstrap for $N = 4$ Super Yang-Mills Amplitudes*, *PoS CORFU2019* (2020) 003 [2005.06735].
- [25] A. Alnuqaydan, S. Gleyzer and H. Prosper, *SYMBA: symbolic computation of squared amplitudes in high energy physics with machine learning*, *Mach. Learn. Sci. Tech.* **4** (2023) 015007 [2206.08901].
- [26] J. Halverson, B. Nelson and F. Ruehle, *Branes with Brains: Exploring String Vacua with Deep Reinforcement Learning*, *JHEP* **06** (2019) 003 [1903.11616].

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

Transforming the Bootstrap

27

- [27] A. Dersy, M. D. Schwartz and X. Zhang, *Simplifying Polylogarithms with Machine Learning*, *Int. J. Data Sci. Math. Sci.* **1** (2024) 135 [2206.04115](#).
- [28] G. Lample and F. Charton, *Deep learning for symbolic mathematics*, in *8th International Conference on Learning Representations*, 2020, [1912.01412](#).
- [29] S. d'Ascoli, P. Kamienny, G. Lample and F. Charton, *Deep symbolic regression for recurrent sequences*, in *39th International Conference on Machine Learning*, 2022, [2201.04600](#).
- [30] F. Charton, *Learning the greatest common divisor: explaining transformer predictions*, in *12th International Conference on Learning Representations*, 2024, [2308.15594](#).
- [31] F. J. Király, L. Theran, R. Tomioka and T. Uno, *The algebraic combinatorial approach for low-rank matrix completion*, *J. Mach. Learn. Res.* **16** (2015) 1391 [1211.4116](#).
- [32] N. Lee, K. Sreenivasan, J. D. Lee, K. Lee and D. Papailiopoulos, *Teaching arithmetic to small transformers*, in *12th International Conference on Learning Representations*, 2024, [2307.03381](#).
- [33] L. Brink, J. H. Schwarz and J. Scherk, *Supersymmetric Yang-Mills Theories*, *Nucl. Phys. B* **121** (1977) 77.
- [34] L. J. Dixon, A. J. McLeod and M. Wilhelm, *A Three-Point Form Factor Through Five Loops*, *JHEP* **04** (2021) 147 [2012.12286](#).
- [35] A. Brandhuber, G. Travaglini and G. Yang, *Analytic two-loop form factors in $N=4$ SYM*, *JHEP* **05** (2012) 082 [1201.4170](#).
- [36] C. Duhr, *Hopf algebras, coproducts and symbols: an application to Higgs boson amplitudes*, *JHEP* **08** (2012) 043 [1203.0454](#).
- [37] A. B. Goncharov, M. Spradlin, C. Vergu and A. Volovich, *Classical Polylogarithms for Amplitudes and Wilson Loops*, *Phys. Rev. Lett.* **105** (2010) 151605 [1006.5703](#).
- [38] S. Golkar, M. Pettee, M. Eickenberg, A. Bietti, M. Cramer, G. Krawezik et al., *xval: A continuous number encoding for large language models*, [2310.02989](#).
- [39] R. F. Nogueira, Z. Jiang and J. Lin, *Investigating the limitations of the transformers with simple arithmetic tasks*, in *1st Mathematical Reasoning in General Artificial Intelligence Workshop, 9th International Conference on Learning Representations*, 2021, [2102.13019](#).
- [40] E. Wenger, M. Chen, F. Charton and K. Lauter, *Salsa: Attacking lattice cryptography with transformers*, in *36th Annual Conference on Neural Information Processing Systems*, 2022, [2207.04785](#).
- [41] T. B. Brown, B. Mann, N. Ryder, M. Subbiah, J. Kaplan, P. Dhariwal et al., *Language models are few-shot learners*, in *34th Annual Conference on Neural Information Processing Systems*, 2020, [2005.14165](#).
- [42] H. Touvron, T. Lavril, G. Izacard, X. Martinet, M.-A. Lachaux, T. Lacroix et al., *Llama: Open and efficient foundation language models*, [2302.13971](#).
- [43] D. P. Kingma and J. Ba, *Adam: A method for stochastic optimization*, in *3rd International Conference on Learning Representations*, 2015, [1412.6980](#).
- [44] A. Paszke, S. Gross, F. Massa, A. Lerer, J. Bradbury, G. Chanan et al., *Pytorch: An imperative style, high-performance deep learning library*, in *33rd Annual Conference on Neural Information Processing Systems*, 2019, [1912.01703](#).
- [45] S. Hochreiter and J. Schmidhuber, *Long short-term memory*, *Neural Computation* **9** (1997) 1735.
- [46] K. Cho, B. van Merriënboer, C. Gulcehre, F. Bougares, H. Schwenk and Y. Bengio, *Learning phrase representations using rnn encoder-decoder for statistical machine translation*, in *Conference on Empirical Methods in Natural Language Processing (EMNLP 2014)*, 2014.
- [47] G. W. Merz, F. Charton, T. Cai, K. Cranmer, L. Dixon, N. Nolte et al., *Transformers for scattering amplitudes*, in *Machine Learning and the Physical Sciences: Workshop at the 37th conference on Neural Information Processing Systems (NeurIPS)*. Non-archival., 2023, https://ml4physicalsciences.github.io/2023/files/NeurIPS_ML4PS_2023_72.pdf.
- [48] A. B. Goncharov, *Multiple polylogarithms, cyclotomy and modular complexes*, *Math. Res. Lett.* **5** (1998) 497 [1105.2076](#).
- [49] A. B. Goncharov, *Multiple polylogarithms and mixed Tate motives*, [math/0103059](#).

Transforming the Bootstrap

28

- [50] F. C. S. Brown and A. Levin, *Multiple Elliptic Polylogarithms*, 1110.6917.
- [51] J. L. Bourjaily et al., *Functions Beyond Multiple Polylogarithms for Precision Collider Physics*, in *Snowmass 2021*, 3, 2022, 2203.07088.
- [52] P. Shaw, J. Uszkoreit and A. Vaswani, *Self-attention with relative position representations*, in *16th Annual Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 2018, 1803.02155.
- [53] Z. Dai, Z. Yang, Y. Yang, J. G. Carbonell, Q. V. Le and R. Salakhutdinov, *Transformer-xl: Attentive language models beyond a fixed-length context*, in *57th Conference of the Association for Computational Linguistics*, 2019, 1901.02860.
- [54] A. Gulati, J. Qin, C.-C. Chiu, N. Parmar, Y. Zhang, J. Yu et al., *Conformer: Convolution-augmented transformer for speech recognition*, in *21st Annual Conference of the International Speech Communication Association*, 2020, 2005.08100.
- [55] J. Su, Y. Lu, S. Pan, B. Wen and Y. Liu, *Roformer: Enhanced transformer with rotary position embedding*, *Neurocomputing* **568** (2024) 127063 2104.09864.
- [56] Y. Sun, L. Dong, B. Patra, S. Ma, S. Huang, A. Benhaim et al., *A length-extrapolatable transformer*, in *61st Annual Meeting of the Association for Computational Linguistics*, 2023, 2212.10554.
- [57] P. Wang, “rotary-embedding-torch.” <https://github.com/lucidrains/rotary-embedding-torch> 2023.
- [58] A. Kazemnejad, I. Padhi, K. N. Ramamurthy, P. Das and S. Reddy, *The impact of positional encoding on length generalization in transformers*, in *37th Annual Conference on Neural Information Processing Systems*, 2023, 2305.19466.

Accepted Manuscript

Appendix A. Mapping Generalized Polylogarithms to Symbols

One of the major difficulties in computing multi-loop amplitudes is the complexity of the multivariate transcendental functions that are encountered. At one loop, besides the logarithm, only one special function appears, namely the classical dilogarithm

$$\text{Li}_2(x) = - \int_0^x \frac{dt}{t} \ln(1-t) = \int_0^1 \frac{dt}{t} \int_0^t \frac{dt'}{1-t'}. \quad (\text{A.1})$$

We define the *weight* of a polylogarithmic function as the number of non-trivial integrations over rational functions it requires. Thus, the classical dilogarithm has weight two, because there are two such integrations. At L loops, functions with weight up to $2L$ are encountered, including generalized polylogarithms [48, 49], defined recursively by

$$G(a_1, a_2, \dots, a_n, x) = \int_0^x \frac{dt}{t - a_1} G(a_2, \dots, a_n, t), \quad G(\vec{0}_n, x) = \frac{1}{n!} \ln^n x, \quad (\text{A.2})$$

with $G = 1$ for $n = 0$. Such functions are quite well understood mathematically, and that understanding has fueled much of the recent progress in multi-loop scattering amplitudes.

In more complex processes, generalized polylogarithms are not enough, and elliptic polylogarithms [50] and functions beyond elliptic are needed; see e.g. ref. [51] and references therein. Fortunately, the three-gluon form factor in SYM can be expressed in terms of the generalized polylogarithms [A.2], with a weight that is precisely $2L$ at loop order L .

Consider a polylogarithmic function P with weight n . It is simplest to describe it iteratively via its derivatives, or total differential, which can be expressed in terms of a number of weight $n - 1$ functions P^{s_k} :

$$dP = \sum_{s_k \in \mathcal{L}} P^{s_k} d \ln s_k, \quad (\text{A.3})$$

where \mathcal{L} is the *symbol alphabet* containing *letters* s_k , which are functions of the underlying kinematical variables.

One can iterate eq. (A.3) n times, defining the weight $n - 2$ functions P^{s_j, s_k} by

$$dP^{s_k} = \sum_{s_j \in \mathcal{L}} P^{s_j, s_k} d \ln s_j, \quad (\text{A.4})$$

and so on, finally reaching the *symbol* [37] of P , $\mathcal{S}[P]$, denoted by

$$\mathcal{S}[P] \equiv \sum_{s_{i_1}, \dots, s_{i_n} \in \mathcal{L}} P^{s_{i_1}, \dots, s_{i_n}} s_{i_1} \otimes \dots \otimes s_{i_n}. \quad (\text{A.5})$$

Here $P^{s_{i_1}, \dots, s_{i_n}}$ is an n -fold tensor of weight-0 polylogarithms, i.e., rational numbers, which we denote by $C^{s_{i_1}, \dots, s_{i_n}}$ in the main text. The \ln 's associated with the $d \ln s_k$'s are by convention omitted from the tensor product defining the symbol.

Transforming the Bootstrap

30

For the three-gluon form factor considered in this article, the symbol alphabet is

$$\mathcal{L}_{3\text{gFF}} = \{a, b, c, d, e, f\}, \quad (\text{A.6})$$

where the six letters read

$$a = \sqrt{\frac{u}{vw}}, \quad b = \sqrt{\frac{v}{wu}}, \quad c = \sqrt{\frac{w}{uv}}, \quad d = \frac{1-u}{u}, \quad e = \frac{1-v}{v}, \quad f = \frac{1-w}{w}, \quad (\text{A.7})$$

and the (u, v, w) are rational functions of the kinematical variables:

$$u = \frac{s_{12}}{q^2}, \quad v = \frac{s_{23}}{q^2}, \quad w = \frac{s_{31}}{q^2} \quad (\text{A.8})$$

with $s_{ij} = (p_i + p_j)^2$. The p_i 's are the relativistic four-momenta associated with the three on-shell external states, i.e., the three gluons, and $q^2 = s_{123} = s_{12} + s_{23} + s_{31}$.

Given the arbitrary nature of ordering the p_i 's, the form factor $\mathcal{F}_{3\text{gFF}}$ is invariant under a dihedral symmetry group D_3 (which is isomorphic to the full permutation group S_3). This group is generated by two transformations which act on the symbol letters as

$$\text{cycle: } \{a, b, c, d, e, f\} \rightarrow \{b, c, a, e, f, d\}, \quad \text{and flip: } \{a, b, c, d, e, f\} \rightarrow \{b, a, c, e, d, f\}. \quad (\text{A.9})$$

Since the symbol involves taking n derivatives, it does not capture all the information in the function. Certain boundary conditions are also required when integrating back up. However, the symbol contains almost all of the critical information, thus serving as a scaffold for determining the full function.

Appendix B. Model Evaluation on the Full List of Linear Relations

In this appendix, we provide a list of linear relations that the form factor symbol satisfies. We also provide results for the accuracy with which the relations hold as a function of training epoch, similar to the discussion in Section 6 but for additional relations.

We recall from Section 2.1 that there are **adjacency** and **prefix/suffix** rules that strongly limit which keys can have nonzero coefficients. Beyond those conditions, there is a large set of linear relationships. We now describe the ones we evaluate during model training. These linear relationships can be grouped into three classes: the triple-adjacency relation, the integrability relations, and the multiple-final-entries relations. We list the relations we consider without providing justification; see [14, 34] for details about the underlying physics.

- **Triple Adjacency:** There is **one** triple adjacency relation (up to dihedral transformations). It relates three terms at a time, and the specified adjacent slots can appear anywhere in the key. When evaluated on the symbol at a certain loop, F can be identified with any of the polylogarithms P described in Appendix

A

$$\text{triple 0: } F^{a,a,b} + F^{a,b,b} + F^{a,c,b} = 0. \quad (\text{B.1})$$

Transforming the Bootstrap

31

- **Integrability:** There are **three** different integrability relations, with the longest one relating 14 terms at a time. Again, the specified adjacent slots can sit anywhere in the key.

$$\text{integ 0: } F^{a,b} + F^{a,c} - F^{b,a} - F^{c,a} = 0, \quad (\text{B.2})$$

$$\text{integ 1: } F^{c,a} + F^{c,b} - F^{a,c} - F^{b,c} = 0, \quad (\text{B.3})$$

$$\begin{aligned} \text{integ 2: } F^{d,b} - F^{d,c} - F^{b,d} + F^{c,d} + F^{e,c} - F^{e,a} - F^{c,e} + F^{a,e} \\ + F^{f,a} - F^{f,b} - F^{a,f} + F^{b,f} + 2F^{c,b} - 2F^{b,c} = 0. \end{aligned} \quad (\text{B.4})$$

- **Multiple Final Entries:** Apart from the trivial zeros, there exist **26** final entry relations involving up to the last four letters. Here the specified adjacent slots can only be at the end of the key, and we highlight this difference by using \mathcal{E} instead of F . When evaluated on the symbol, \mathcal{E} is the same as $\mathcal{F}^{(L)}$.

$$\text{final 0: } \mathcal{E}^{a,d} = 0, \quad \text{final 1: } \mathcal{E}^{e,d} = 0,$$

$$\text{final 2: } \mathcal{E}^{a,d,d} = 0, \quad \text{final 3: } \mathcal{E}^{a,b,d} = 0,$$

$$\text{final 4: } \mathcal{E}^{a,c,e} = 0, \quad \text{final 5: } \mathcal{E}^{e,b,d} = 0, \quad \text{final 6: } \mathcal{E}^{e,d,d} = 0,$$

$$\text{final 7: } \mathcal{E}^{a,d,d,d} = 0, \quad \text{final 8: } \mathcal{E}^{a,b,b,d} = 0, \quad \text{final 9: } \mathcal{E}^{a,d,b,d} = 0,$$

$$\text{final 10: } \mathcal{E}^{c,b,b,d} = 0, \quad \text{final 11: } \mathcal{E}^{e,b,b,d} = 0, \quad \text{final 12: } \mathcal{E}^{e,b,d,d} = 0,$$

$$\text{final 13: } \mathcal{E}^{e,d,b,d} = 0, \quad \text{final 14: } \mathcal{E}^{e,d,d,d} = 0, \quad \text{final 15: } \mathcal{E}^{f,d,b,d} = 0,$$

$$\text{final 16: } \mathcal{E}^{b,f} - \mathcal{E}^{b,d} = 0, \quad (\text{B.5})$$

$$\text{final 17: } \mathcal{E}^{c,d,d} + \mathcal{E}^{c,e,e} = 0, \quad (\text{B.6})$$

$$\text{final 18: } \mathcal{E}^{d,d,b,d} - \mathcal{E}^{d,b,d,d} = 0, \quad (\text{B.7})$$

$$\text{final 19: } \mathcal{E}^{c,b,d,d} - \mathcal{E}^{c,d,b,d} = 0, \quad (\text{B.8})$$

$$\text{final 20: } \mathcal{E}^{f,b,d} - \mathcal{E}^{d,b,d} + \mathcal{E}^{b,d,d} = 0, \quad (\text{B.9})$$

$$\begin{aligned} \text{final 21: } \mathcal{E}^{b,d,d,d} + \mathcal{E}^{f,a,f,f} - \mathcal{E}^{d,b,d,d} - \mathcal{E}^{e,a,f,f} \\ + \mathcal{E}^{f,b,d,d} - \mathcal{E}^{a,e,e,e} = 0, \end{aligned} \quad (\text{B.10})$$

$$\begin{aligned} \text{final 22: } \mathcal{E}^{a,b,d,d} - \frac{1}{2}\mathcal{E}^{c,d,d,d} - \frac{1}{2}\mathcal{E}^{d,c,e,e} + \frac{1}{2}\mathcal{E}^{a,e,e,e} \\ + \frac{1}{2}\mathcal{E}^{e,a,f,f} - \frac{1}{2}\mathcal{E}^{f,a,f,f} + \frac{1}{2}\mathcal{E}^{e,c,e,e} = 0, \end{aligned} \quad (\text{B.11})$$

$$\begin{aligned} \text{final 23: } \mathcal{E}^{c,b,d,d} - \frac{1}{2}\mathcal{E}^{b,f,f,f} + \frac{1}{2}\mathcal{E}^{d,c,e,e} - \frac{1}{2}\mathcal{E}^{e,c,e,e} \\ + \frac{1}{2}\mathcal{E}^{c,d,d,d} + \frac{1}{2}\mathcal{E}^{d,b,d,d} - \frac{1}{2}\mathcal{E}^{f,b,d,d} = 0, \end{aligned} \quad (\text{B.12})$$

$$\begin{aligned} \text{final 24: } \mathcal{E}^{c,d,b,d} - \frac{1}{2}\mathcal{E}^{b,f,f,f} + \frac{1}{2}\mathcal{E}^{d,c,e,e} - \frac{1}{2}\mathcal{E}^{e,c,e,e} \\ + \frac{1}{2}\mathcal{E}^{c,d,d,d} + \frac{1}{2}\mathcal{E}^{d,b,d,d} - \frac{1}{2}\mathcal{E}^{f,b,d,d} = 0, \end{aligned} \quad (\text{B.13})$$

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

Transforming the Bootstrap

32

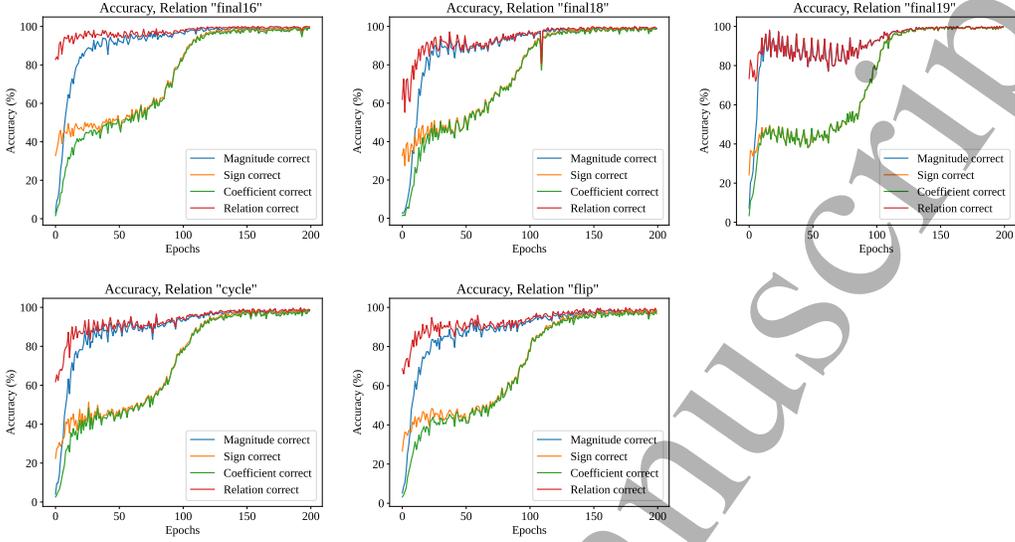


Figure B1. Coefficient accuracy (green), magnitude accuracy (blue), relation accuracy (green), and sign accuracy (yellow) for all Group 1 relations.

$$\begin{aligned}
 \text{final 25: } \mathcal{E}^{f,b,b,d} - \mathcal{E}^{d,b,b,d} + \mathcal{E}^{b,b,d,d} - \frac{1}{2}\mathcal{E}^{f,a,f,f} + \frac{1}{2}\mathcal{E}^{d,b,d,d} \\
 - \frac{1}{2}\mathcal{E}^{f,b,d,d} + \frac{1}{2}\mathcal{E}^{e,a,f,f} + \frac{1}{2}\mathcal{E}^{a,e,e,e} - \frac{1}{2}\mathcal{E}^{b,f,f,f} = 0.
 \end{aligned}
 \tag{B.14}$$

The first sixteen relations (final 0 to final 15) are all one-term relations, restricting which letters *cannot* appear at the end of any key. They are trivial for the model to learn and can be predicted perfectly after just a few epochs. We therefore omit a detailed assessment of these one-term relations.

- **Dihedral Symmetry:** The cycle and flip symmetries can both be expressed as two-term Group 1 relations because they require two keys related by the given symmetry to have identical coefficients. In order to evaluate the cycle and flip relations independently of each other, we select only one two-term relation instance of either kind from each 6-term dihedral orbit.

We evaluate the models on the list of relations above using the same architecture, dataset, and accuracy criteria as in Section 6. We again categorize the relations according to the three groups mentioned in the main text. We plot the Group 1 results in Fig. B1, the Group 2 results in Fig. B2, and the Group 3 results in Fig. B3. We note that Group 3 relations behave somewhat differently, depending on whether they are short (less than 4 terms) or long (more than 6 terms).

Transforming the Bootstrap

33

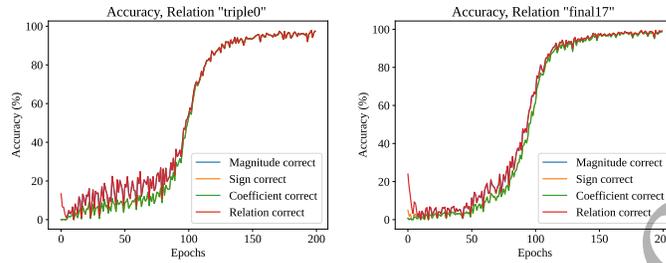


Figure B2. Coefficient accuracy (green), magnitude accuracy (blue), relation accuracy (green), and sign accuracy (yellow) for all Group 2 relations.

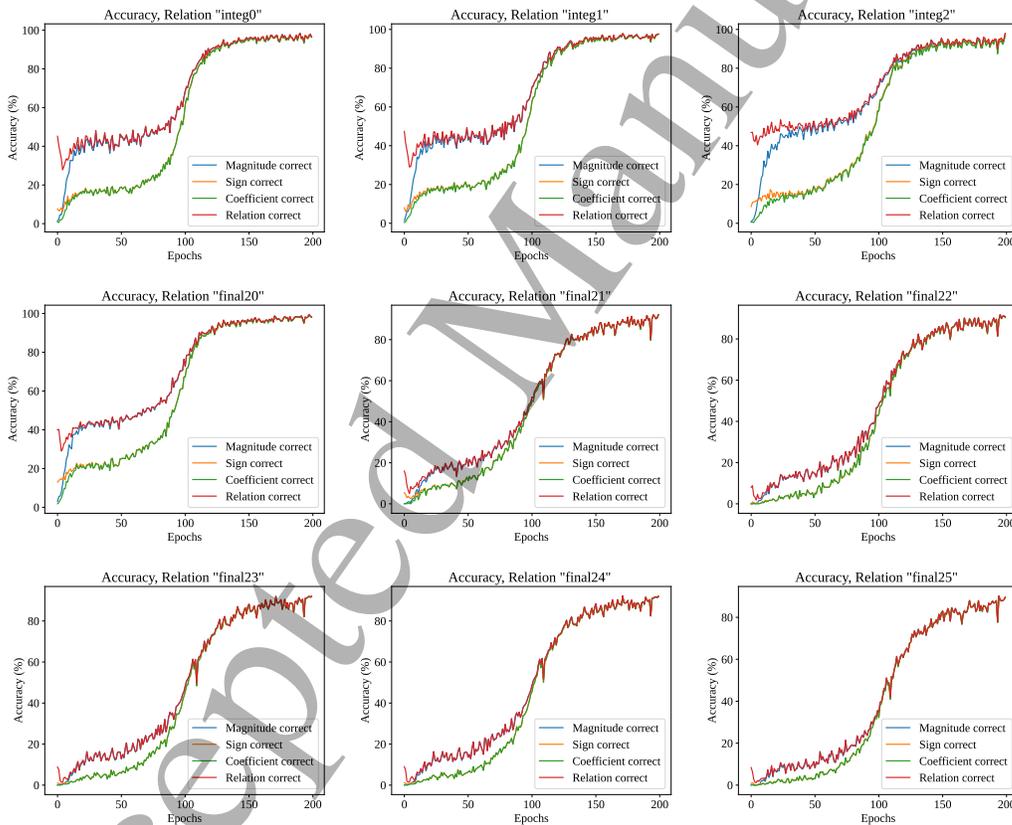


Figure B3. Coefficient accuracy (green), magnitude accuracy (blue), relation accuracy (green), and sign accuracy (yellow) for all Group 3 relations.

Appendix C. Coefficient-from-Key: Model Ablations

We perform a number of ablations in order to better understand the model performance for the coefficient-from-key experiment in Section 4, with an eye towards building a future model that can generalize across loops. Here we consider several different positional

Transforming the Bootstrap

34

encoding schemes, explore encoder-only and decoder-only architectures, and vary both the number of layers in the model and the dimension d of the encoded tokens. See Tables C1 and C3 for an overview of the results.

All models exhibit the two-phase learning structure, suggesting that this behavior is an intrinsic property of the problem itself rather than an artifact of the chosen model architecture. We report the epoch at which the test-set magnitude accuracy first exceeds 95%, the midpoint epoch of the plateau step (defined as the epoch at which the model’s overall accuracy first reaches the midpoint between its final-state accuracy and 75%), the epoch at which the overall test-set accuracy first exceeds 95%, and the best overall test-set accuracy after 200 epochs.

Appendix C.1. Architecture Ablations: Relative Position Encoding, Encoder-only, Decoder-only, and Non-Transformer Architectures

First, we test whether results change appreciably when the sign token is positioned at the end of the coefficient rather than the beginning, leaving all other hyperparameters unchanged from the baseline configuration. We next explore three variants of relative position encoding in the encoder of the baseline model of Section 4: (1) one in which learned relative position encoding is added to the product of the key and query matrices [52] and implemented as in the TransformerXL [53] and Conformer [54] family of models; and (2-3) two rotary encodings—one standard (Rotary Position Embedding, a.k.a. RoPE [55]) and one long-context (xPos) [56]. RoPE implements relative position encoding by multiplying the query and key matrices by a rotation matrix, while xPos introduces an additional scaling factor to enable better attention to tokens separated by large distances. Both are implemented in a common repository [57]. We compare to a baseline four-layer encoder-decoder Transformer with $d = 256$. In all relative-position-encoding experiments, we only apply positional encoding in the encoder and do not use positional encoding at all in the autoregressive decoder (as in NoPE [58]). For these experiments, we use the same zero percentage and nontrivial-zero biased prescription as in the model described in Section 6.

We find that models trained with relative position encoding slightly lag behind models trained with absolute position encoding, but achieve comparable test-set accuracy after the full 200 epochs. Additionally, we note that differences between RoPE and xPos are small.

Additionally, for the baseline model, varying the random seed three times can cause the point at which the test-set magnitude accuracy first exceeds 95% (the first column of Table to shift by up to 5 epochs and the point at which the test-set magnitude accuracy first exceeds 95% to shift by up to 8 epochs— thus, many of the differences in performance between architectures are within this window of uncertainty.

Our second class of ablations explores different Transformer architectures. In addition to the encoder-decoder Transformer used for the majority of our experiments, we train both an encoder-only Transformer and a decoder-only Transformer and evaluate

their performance. The encoder-only (or decoder-only) model has 8 layers in the encoder (or decoder), $d = 256$, and 8 attention heads. We also train an encoder-only Transformer with 8 layers, $d = 256$, and 8 attention heads using the RoPE rotary positional embedding.

In order to perform training in a decoder-only paradigm, we concatenate coefficient and key in a “prompt-completion” manner similar to that of the GPT family of models [41]. When evaluating the predictions of the decoder-only model, we only evaluate accuracy on the “completion” portion, which is the tokens corresponding to the integer coefficients. Although the model is trained on next-token prediction on the full input, whether it correctly outputs the next letter in the key is not a well-defined objective, as many possible keys exist. A single solution exists only once the model has seen all tokens corresponding to the full key and is asked to predict the remaining tokens corresponding to the coefficient.

Likewise, when evaluating the predictions of the encoder-only model, all model outputs will be exactly $2L$ tokens long (the same length as the input). Because the coefficient sequence encoded in base-1000 encoding will be much shorter than $2L$, the encoder-only model will produce more output tokens than there are tokens in the coefficient sequence. We require that the tokens at the start of this output sequence correspond to the tokens of the coefficient sequence and that the output token immediately after the coefficient sequence corresponds to an end-of-sequence token. We ignore all tokens in the encoder output that appear after this end-of-sequence token.

Architecture	First Mag. Acc. > 95% [Epoch]	Midpoint of Step [Epoch]	First Total Acc. > 95% [Epoch]	Best Acc. Epoch 200
4/4 enc-dec	26	94	105	99.2%
4/4 enc-dec, sign last	27	102	114	99.5%
4/4 enc-dec, rel_pos	34	98	105	98.4%
4/4 enc-dec, RoPE	33	121	132	99.1%
4/4 enc-dec, xPos	38	121	133	99.3%
8-layer, enc-only	28	107	115	99.5%
8-layer, enc-only, RoPE	36	122	131	99.4%
8-layer, dec-only	39	105	112	99.2%
8-layer, dec-only, sign last	39	118	134	99.3%

Table C1. Model learning dynamics for a variety of architectures trained and evaluated on the same $L = 6$ dataset. All display the double-plateau structure. We indicate the epoch at which the test-set magnitude accuracy first exceeds 95%, the midpoint of the plateau step, the epoch at which the overall test-set accuracy first exceeds 95%, and the best overall test-set accuracy after 200 epochs. Numbers reported are for the best result of three seeds, with ‘best’ defined as the first model to reach 95% accuracy, or, if this threshold is not reached by any of the three, the model with the highest overall test-set accuracy after 200 epochs. The baseline model is indicated in bold.

Finally, we perform the nonzero coefficient-from-key prediction task at 5 loops using non-Transformer architectures. Specifically, we use an encoder-decoder Long-Short-Term Memory recurrent neural network (LSTM) [45] and a Gated Recurrent Unit (GRU) [46].

Transforming the Bootstrap

36

The embedding dimension is $d = 512$ and the number of layers in the encoder and decoder are both taken to be 2, as in the Transformer, while the hidden layer size is either chosen such that the number of trainable parameters is approximately the same as in the Transformer, or chosen to be $h = 2048$. We again observe the familiar two-phase behavior as in Fig. 3. We report the results compared to a Transformer model in the table C2. For the models with a comparable number of parameters as the Transformer, we observe that performance is slightly worse. However, when $h = 2048$, the LSTM performs at least as well as the Transformer on this particular task, and the GRU is not qualitatively worse.

Architecture	First Mag. Acc. > 95% [Epoch]	Midpoint of Step [Epoch]	First Total Acc. > 95% [Epoch]	Best Acc. Epoch 15	N_{params}
Transformer	3	6	9	99.10%	20M
LSTM ($h = 786$)	7	12	N/A	87.42%	20M
GRU ($h = 916$)	5	10	N/A	90.91%	20M
GRU ($h = 2048$)	3	5	6	98.88%	85M
LSTM ($h = 2048$)	3	5	6	99.76%	112M

Table C2. Model learning dynamics for a variety of LSTM architectures trained and evaluated on the same $L = 5$ dataset. All display the double-plateau structure. We indicate the epoch at which the test-set magnitude accuracy first exceeds 95%, the midpoint of the plateau step, the epoch at which the overall test-set accuracy first exceeds 95%, and the best overall test-set accuracy after 15 epochs. The baseline Transformer model is indicated in bold. For this experiment, we use only one seed rather than the best of 4.

Appendix C.2. Architecture Parameter Scan

We now scan a variety of model depths and sizes for the $L = 6$ coefficient-from-key prediction task, with results given in Table C3. We vary the number of layers in both the encoder and decoder and the internal dimension d of the Transformer. For each model, we report the number of epochs required to reach a variety of benchmarks that summarize the two-phase learning behavior: (1) the number of epochs required to reach the first phase plateau (> 95% magnitude accuracy); (2) the number of epochs required to reach the midpoint of the step between the two plateaus (i.e., the midpoint between 75% overall accuracy and the value of overall accuracy after 200 epochs); and (3) the number of epochs required to reach 95% overall accuracy. We also report the best overall accuracy after 200 epochs.

Qualitatively, while all models exhibit the two-phase learning behavior, both phases are compressed (last for fewer epochs) as model capacity increases. The relationship between the number of training epochs required to reach a particular benchmark and the model dimension appears to be approximately linear. That is, increasing model dimension by a factor of two while leaving all other hyperparameters fixed will cause benchmarks to be reachable in half as many epochs. The number of attention heads

does not appear to have as much impact on learning dynamics as the model dimension and depth.

Architecture	First Mag. Acc. > 95% [Epoch]	Midpoint of Step [Epoch]	First Total Acc. > 95% [Epoch]	Best Acc. Epoch 200
4/4, $d = 256, h = 8$	26	94	105	99.2%
4/4, $d = 256, h = 4$	29	100	108	99.1%
4/4, $d = 512, h = 8$	15	54	62	99.4%
4/4, $d = 512, h = 4$	19	65	75	99.3%
2/2, $d = 256, h = 8$	45	186	N/A	91.7%
2/2, $d = 256, h = 4$	63	N/A	N/A	75.5%
2/2, $d = 512, h = 8$	26	93	107	98.9%
2/2, $d = 512, h = 4$	29	120	133	98.3%

Table C3. Model learning dynamics for a variety of model sizes trained and evaluated on the same $L = 6$ dataset. All display the double-plateau structure. We indicate the epoch at which the test-set magnitude accuracy first exceeds 95%, the midpoint of the plateau step, the epoch at which the overall test-set accuracy first exceeds 95%, and the best overall test-set accuracy after 200 epochs. Numbers reported are for the best result of three seeds, with ‘best’ defined as the first model to reach 95% accuracy, or, if this threshold is not reached by any of the three, the model with the highest overall test-set accuracy after 200 epochs. The baseline model is indicated in bold.

Additionally, because we only split our data into training and test sets, rather than training, validation, and test sets, the accuracy metrics quoted may be unduly influenced by the specific choice of test set. In order to demonstrate that performance does not degrade appreciably for different choices of test set, we perform the 5-loop nonzero coefficient-from-key prediction task with three different random training/test set configurations, but the same architecture and hyperparameter configuration as in the baseline experiment in Section 4. The accuracies for the different splits are a little lower than for the baseline.

Architecture	First Mag. Acc. > 95% [Epoch]	Midpoint of Step [Epoch]	First Total Acc. > 95% [Epoch]	Best Acc. Epoch 15
Baseline	3	6	9	99.10%
Split 1	2	6	9	98.61%
Split 2	3	7	9	98.40%
Split 3	3	7	11	98.92%

Table C4. Model learning dynamics for four different choices of random training and test set split on the 5-loop coefficient-from-key task. We indicate the epoch at which the test-set magnitude accuracy first exceeds 95%, the midpoint of the plateau step, the epoch at which the overall test-set accuracy first exceeds 95%, and the best overall test-set accuracy after 15 epochs.

Appendix D. Further Strike-out Experiments

Table D1 presents the overall, magnitude and sign accuracy of models trained for up to 700 epochs, for different variations of the strike-two method, performed on loop $L = 6$ data as in the main text. We use 4-layer Transformers with dimension $d = 512$ and 8 attention heads. All models are trained for up to 700 epochs, and we indicate the epoch at which the indicated best accuracy is first achieved. The experiments for the “strike-two parents”, “shuffled parents”, “parent signs only” and “parent magnitudes only” sections of Table D1 are described in Section 8; Table D1 provides more information about their dependence on the strike-out distance k .

Here, two new sets of experiments are introduced. In the “sorted unique” experiments, all parents are sorted in ascending order, and duplicate parent coefficients are removed; thus, the multiplicity of each of the parent coefficients, a permutation-invariant property, is eliminated. This model proves harder to train, with the best models achieving an accuracy of 83.4% after 350 epochs. Still, it is much better than random guessing, with sign information almost totally recoverable (93.6%). This result suggests that some amount of information useful to reconstructing the coefficient is present even when multiplicities are removed.

In the “zero/nonzero” experiments, all nonzero parent coefficients are encoded as “1”, while all zero parent coefficients are encoded as “0”. Such experiments provide further information on the “signs-only” experiment in the main text. While clarifying which parents are positive, negative, and zero is enough to reconstruct the sign and magnitude of the target coefficient at 99.0% and 93.5% respectively, providing only which parents are zero or nonzero causes magnitude accuracy to drop to 60.1% and sign accuracy to drop to 61.7%. Thus, zero/nonzero information alone is not enough to reconstruct the target coefficient.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60

Transforming the Bootstrap

39

	Distance	Best Epoch	Train Size	Accuracy	Magnitude Accuracy	Sign Accuracy
Strike-two parents	Full	455	757,500	98.1%	98.4%	99.6%
	5	524	754,060	98.3%	98.6%	99.7%
	3	601	738,352	98.4%	98.7%	99.7%
	2	681	688,869	98.1%	98.3%	99.5%
	1	646	576,510	94.3%	95.2%	98.5%
Shuffled parents	Full	407	4,906,466	95.2%	99.1%	96.3%
	5	376	4,906,466	94.7%	98.8%	95.8%
	3	408	4,906,436	95.1%	99.0%	96.1%
	2	433	4,906,249	93.5%	98.1%	95.0%
	1	442	4,882,501	91.1%	92.2%	96.5%
Sorted parents	Full	389	591,864	91.5%	93.8%	96.6%
	5	432	717,534	93.9%	95.4%	97.9%
	3	514	702,363	93.1%	94.6%	97.5%
	2	453	657,863	90.7%	92.4%	96.5%
	1	459	536,588	76.8%	79.3%	90.5%
Sorted unique parents	Full	329	476,932	79.0%	84.0%	92.1%
	5	355	538,325	83.4%	87.4%	93.6%
	3	349	487,813	79.7%	84.1%	91.7%
	2	287	436,012	73.6%	78.5%	89.4%
	1	314	355,147	57.2%	61.9%	80.9%
Zero / nonzero	Full	304	497,112	40.7%	60.1%	61.7%
	5	229	467,871	35.0%	53.8%	59.4%
	3	93	415,230	22.9%	39.1%	54.9%
	2	18	344,831	10.1%	20.1%	50.3%
	1	1	131,812	0.8%	1.4%	50.4%
Parent signs only	Full	404	748,088	93.3%	93.5%	99.0%
	5	330	739,479	92.4%	92.5%	99.0%
	3	294	711,450	88.2%	88.4%	98.2%
	2	331	653,368	73.3%	73.6%	94.8%
	1	22	468,339	5.8%	6.5%	64.6%
Parent magnitudes only	Full	395	751,675	81.8%	98.4%	83.2%
	5	445	747,187	81.2%	98.4%	82.4%
	3	452	726,554	79.7%	98.3%	80.9%
	2	611	672,561	77.8%	97.9%	79.4%
	1	509	527,843	64.1%	93.6%	67.7%

Table D1. Overall, magnitude and sign accuracy for strike-two method. We show the best of four seeds, and for a variety of experiments and k -values we quote the epoch at which the best overall accuracy was attained, the training set size after duplicate removal, and the best values of overall, magnitude and sign accuracy after 700 epochs.

1
2
3
4
5
6
7
8
9
10
11
12
13
14
15
16
17
18
19
20
21
22
23
24
25
26
27
28
29
30
31
32
33
34
35
36
37
38
39
40
41
42
43
44
45
46
47
48
49
50
51
52
53
54
55
56
57
58
59
60