

Towards Autoformalization of Mathematics and Code Correctness: Experiments with Elementary Proofs

Garett Cunningham
School of EECS
Ohio University
Athens, OH 45701
gc974517@ohio.edu

Razvan C. Bunescu
Department of Computer Science
UNC Charlotte
Charlotte, NC 28223
razvan.bunescu@uncc.edu

David Juedes
School of EECS
Ohio University
Athens, OH 45701
juedes@ohio.edu

Abstract

The ever-growing complexity of mathematical proofs makes their manual verification by mathematicians very cognitively demanding. Autoformalization seeks to address this by translating proofs written in natural language into a formal representation that is computer-verifiable via interactive theorem provers. In this paper, we introduce a semantic parsing approach, based on the Universal Transformer architecture, that translates elementary mathematical proofs into an equivalent formalization in the language of the Coq interactive theorem prover. The same architecture is also trained to translate simple imperative code decorated with Hoare triples into formally verifiable proofs of correctness in Coq. Experiments on a limited domain of artificial and human-written proofs show that the models generalize well to intermediate lengths not seen during training and variations in natural language.

1 Introduction

To the uninitiated, the notion of mathematical proof represents simply an argument written by people to convince others of mathematical truth. However, in a real sense, mathematical proof must have formal underpinnings that go beyond the written argument. Arguments that lack such underpinnings might have fatal errors or even logical inconsistencies (see, for example, Russell’s Paradox (Irvine and Deutsch, 2021)). Nevertheless, mathematical arguments written in natural language are the norm and they have great value.

In Tymoczko (1979)’s well-known paper that discusses a somewhat controversial (at the time) proof of the Four Color Theorem (Appel and Haken, 1977; Appel et al., 1977), he explores “what is a mathematical proof?” He posits that all mathematical proofs must be (i) convincing, (ii) surveyable, and (iii) formalizable. The first two points are for the reader—proofs must be convincing to and comprehensible by mathematicians. For the third point,

he notes that, “Most mathematicians and philosophers believe that any acceptable proof can be formalized. We can always find an appropriate formal language and theory in which the informal proof can be embedded and ‘filled out’ into a rigorous formal proof.” For most mathematicians, this third part is crucial for ensuring that subtle, but fatal, errors in logic do not exist in mathematical proof.

Great progress has been made since the 1970’s in fully formalizing significant mathematical results. For instance, the Feit-Thompson Theorem (Gonthier et al., 2013; Gonthier, 2013) and the Four Color Theorem (Gonthier, 2008) have been formally verified using the proof assistant Coq (Bertot and Castéran, 2013), and the Kepler Conjecture (Hales, 2005; Hales et al., 2017) has been formally verified using the proof assistants Isabelle and HOL Light (Nipkow et al., 2002). Moreover, proof assistants have demonstrated immense utility for software verification, such as the full certification of a C compiler (Leroy, 2009). Proofs demonstrating the correct behavior of code share a similar structure to proofs in pure mathematics, where systems like Hoare logic replace standard first-order logic. Thus, Tymoczko’s criteria for mathematical proof can be extended to the verification of programs. For many experts, LaTeX provides an excellent tool for satisfying the first two criteria. In addition, carefully written LaTeX (Higham, 2020) provides a rich structure for establishing the third criterion.

The vast majority of modern mathematics is expressed using natural language (NL), with the overwhelming majority typeset in LaTeX. Fully formalizing mathematics using proof assistants is still a difficult and time consuming task. This paper takes some preliminary steps toward bridging this gap by exploring how modern machine learning techniques can be used to convert carefully written LaTeX into equivalent, and formally verified mathematics in Coq, a process referred to as *autoformalization* in the literature (Szegedy, 2020).

Wang et al. (2018, 2020) explored the similar task of translating mathematical statements from LaTeX into Mizar, using LSTM-based models with attention. To generate aligned LaTeX-Mizar pairs, they use a tool (Bancerek, 2006) that translates top-level Mizar statements into artificial LaTeX sentences, a task that is facilitated by the fact that Mizar is human readable and similar in length with the corresponding LaTeX version. Carman (2021) evaluated the competency of LSTMs toward formalizing a restricted set of artificially generated theorems about simple arithmetic expressions, reporting reasonable success over expression lengths seen during training. More recently, Wu et al. (2022) evaluated Codex and PaLM on a significantly more limited, but human-written set of theorems in algebra and number theory.

In contrast to prior work, we address the autoformalization of both *theorems and their proofs*, and extend the scope to *proofs of code correctness*. We use a number of manually written mathematical statements to abstract a complex grammar that is then used to generate a dataset of substantially longer and more diverse mathematical theorems and proofs. We develop an architecture based on the Universal Transformer (Dehghani et al., 2018) and adapt a copying mechanism (Gu et al., 2016) to handle arbitrary numbers and variable names at test time. The models are evaluated extensively on their ability to systematically generalize to statement lengths not seen during training, for which we report *sequence-level* accuracy as well as a *semantic-level* accuracy calculated by combining sequence-level accuracy for the theorem and running Coq to determine if the generated proof is correct. Code and data are made available at <https://github.com/gc974517/autoformalization>.

2 Dataset of Theorems and Proofs

We create two independent datasets of mathematical statements that overall correspond to four classes of theorems and proofs: the first dataset contains three classes of arithmetic statements (EVEN-ODD, COMPOSITES, and POWERS), described in detail in Section 2.1, and the second dataset containing statements about code correctness via Hoare logic (POLY), described in detail in Section 2.2. In each example, the input theorem-proof pair is given in LaTeX, whereas the formalized output is represented in Coq. This work focuses on the proof assistant Coq (Bertot and Castéran, 2013) because

(a) there is a rich set of mathematical libraries that have been developed for it, (b) it has been used successfully to reason about significant computation artifacts, such as the ComperCert C compiler (Leroy, 2009), and (c) it benefits from a rich set of training material for the proof assistant related to software verification (Pierce et al., 2010).

Each class of examples demonstrates features necessary for the successful autoformalization of mathematical theorems and proofs. For example, POWERS and COMPOSITES examples may define useful terminology to make the theorems shorter, e.g. proving that 4 is a *square*, or conversely they may state theorems directly without any preliminary definitions, e.g. proving $\exists n. n^2 = 4$. As shown in Figures 3 and 4, this corresponds in Coq to aliasing propositions using the `Definition` keyword. Additionally, the examples in the dataset provide a stress test of the copying mechanism described in Section 3.1, testing its ability to learn the correct order and number of terms to include in mathematical expressions, as well as their placement in theorems and proofs, in a way that generalizes to arbitrary tokens in mathematical language.

For each of the four classes of theorems and proofs, we manually created a few examples ourselves in order to guide the construction of a complex grammar that is then used to generate a dataset of substantially longer and more diverse mathematical theorems and proofs. Each dataset is generated using its corresponding grammar in an identical way. First, a random seed is sampled that controls the overall structure of the theorem, proof, and definition, if any. Then, the skeleton structure of the proof is completed with phrases that are sampled from a separate context-free grammar. The coarse control of the skeleton structure allows the construction of examples with interesting features like sublemmas, forward or backward proof direction, coreference, or additional conditions for the theorem, among others.

Many of the difficulties in formalizing mathematical statements from NL into Coq stem from the wide variability in the level of detail of mathematical proofs, and the frequent mismatch between what is considered an acceptable inference step in NL proofs vs. an inference step in Coq. Furthermore, there may be multiple Coq proofs for any given theorem, at different levels of granularity. We address this ambiguity by requiring the structure of the Coq proof to match the overall structure of the

Theorem. $28M + 308$ is even.

Proof. We know the summation between even numbers in \mathbb{N} will be an even number. Observe that 308 is known to be even. Additionally, note that the pair $M \times 28$ is trivially even. This is true because the coefficient 28 is even. \square

Require Import Arith.

Theorem M28_308:

forall M : nat, Nat.even(28 * M + 308) = true.

Proof.

```
intros.
repeat rewrite Nat.even_add.
assert (H1: Nat.even 308 = true).
{ auto. }
assert (H2: Nat.even 28 = true).
{ auto. }
assert (H3: Nat.even (28 * M) = true).
{ rewrite Nat.even_mul.
  rewrite H2.
  auto. }
rewrite H1.
rewrite H3.
auto.
```

Qed.

Figure 1: Generated example from the EVEN-ODD set.

NL proof. This is achieved by quasi-synchronously generating the LaTeX and Coq versions of mathematical statements, while still allowing for some simple re-orderings in order to improve generalization performance, e.g. swapping arguments of commutative operations.

In total, the grammar-based method for generating examples can theoretically produce over 283 million unique arithmetic examples and over 491,000 unique code examples, before considering variations in phrasing by sampling from the context-free grammar.

2.1 Arithmetic Statements

We generated three classes of mathematical statements, i.e. theorem-proof pairs:

- EVEN-ODD: an expression is even or odd.
- COMPOSITES: a number is composite.
- POWERS: a number is an integer power of n .

EVEN-ODD examples contain arithmetic expressions of n variables with even coefficients that are summed with a constant term, meaning that the parity of this constant determines the parity of the whole expression. Proofs make use of this fact with varying rigor based on our manually designed

Theorem. $450 + a \cdot 192 + j \cdot 462$ is guaranteed to be even for any natural terms j , and a .

Proof. It can be justified that $192 \cdot a + j \cdot 462$ is trivially even. Note that $192a$ is an even number in \mathbb{N} because multiplying between an even integer with an arbitrary number in \mathbb{N} is guaranteed to be even. Likewise, $462j$ is trivially an even number in \mathbb{N} . The claim is proven as a consequence of the fact that the sum of even numbers with an even number will be in itself an even number. Therefore, our theorem holds. \square

Require Import Arith.

Theorem a450_192j_450_even:

forall j a : nat,
Nat.even (192 * a + 462 * j + 450) = true.

Proof.

```
intros.
rewrite Nat.even_add.
assert (H1: Nat.even (192 * a) = true).
{ rewrite Nat.even_mul.
  auto. }
assert (H2: Nat.even (462 * j) = true).
{ rewrite Nat.even_mul.
  auto. }
assert (H3: Nat.even
(192 * a + 462 * j) = true).
{ repeat rewrite Nat.even_add.
  rewrite H1.
  rewrite H2.
  auto. }
rewrite H3.
auto.
```

Qed.

Figure 2: Instance of sublemma use in the EVEN-ODD dataset. The proof that the sum of non-constant terms is even (assertion H3) is given before proving the theorem.

grammar, an example of which is shown by Figure 1. The Coq program is generated concurrently with the paired LaTeX example. The example shown in Figure 2 illustrates the generation and use of prior facts to prove an implicit sublemma, in both the natural language and matching Coq version.

Examples of theorems and proofs for POWERS and COMPOSITES share a similar structure in both their LaTeX and Coq forms, as shown in Figures 3 and 4, respectively. The theorems assert the existence of a natural number such that a defining property holds and their proofs are constructive, with the distinction that examples for composites prove factorization into n factors.

For both training and testing, we generate 5,000 even-odd, 5,000 composites, and 2,000 powers examples. We train on values of $n \in \{2, 3, 5, 7, 9\}$ and test on values $n \in \{2, 3, \dots, 12\}$, where n rep-

Definition. We define that $w \in \mathbb{N}$ is a composite natural number if taking some $R, Q \in \mathbb{N}$ we have $Q, R \geq 2$ and $Q \times R = w$.

Theorem. 35 is a composite whole number.

Proof. Remember that a composite natural number is the multiplication between Q and R such that Q and $R \geq 2$. Allow $R = 7, Q = 5$. We justify the result is valid as $35 = Q \cdot R$. \square

```
Require Import Lia.
Definition composite
  (w : nat) :=
    exists R Q : nat,
      (Q >= 2) /\ (R >= 2)
      /\ (Q * R = w).
Theorem w_composite:
  composite 35.
Proof.
  unfold composite.
  exists 7.
  exists 5.
  lia.
Qed.
```

Figure 3: Generated COMPOSITES example.

Definition. We say o is a square whenever there exists some whole number Z such that $Z \geq 2$ and $o = Z^2$.

Theorem. $o = 64$ is a square.

Proof. Let $Z = 8$. Observe that $64 = 8^2$. Also notice $Z = 8$ is more than or equal to 2. This yields 64 is a square whole number. \square

```
Require Import Lia.
Definition square
  (o : nat) :=
    exists Z : nat,
      (Z >= 2) /\ (o = Z^2).
Theorem square_64:
  square 64.
Proof.
  unfold square.
  exists 8.
  assert (H1: 8 >= 2).
  { lia. }
  repeat split.
  apply H1.
Qed.
```

Figure 4: Generated example from the POWERS set.

resents the number of variables in the arithmetic expression, the number of factors, or the power, respectively. This is done in order to evaluate the model’s ability to generalize to unseen arithmetic

expression lengths and numbers of factors.

2.1.1 Handwritten Examples

We also created a small collection of 45 human-written LaTeX theorem-proof pairs to evaluate performance on examples outside of our manually generated grammar. These are distinct from the original manually written examples that were used to guide the development of the generative grammar. There are 15 examples for each type of proof from the arithmetic set, using the same vocabulary with a number of unseen grammatical structures.

2.2 Code Correctness Statements

We create a dataset of correctness proofs about short programs written in the imperative programming language *Imp* (Pierce et al., 2018), which we call POLY. The programs represent various algorithms for evaluating a polynomial, and their proofs of correctness verify that the programs correctly model the polynomial as a mathematical function. Proofs are conducted as either fully decorated programs or as sequences of Hoare triples with natural language justifying steps in between. An example is shown in Figure 5.

For both training and testing data, we generate 5,000 examples. We train on programs containing 2, 3, 5, 7, 9, and 11 lines, then test on programs containing from 2 up to 14 lines to evaluate the model’s ability to generalize to novel program lengths.

3 Semantic Parsing Architecture

To formalize LaTeX statements into Coq, we developed an encoder-decoder architecture based on the Universal Transformer (Dehghani et al., 2018). Similar to Csordás et al. (2021), we do so by adding recursive passes into the encoder and decoder of a base Transformer (Vaswani et al., 2017), thus making the model analogous to a Universal Transformer without adaptive computation time (ACT). Further, we introduce a copying mechanism and support for out-of-vocabulary mathematical terms.

3.1 Copying Mechanism

Mathematical language contains features uncommon or non-existent in natural language, such as numbers, variables, and carefully defined terminology. In order to address the use of general mathematical jargon, these tokens are replaced in the LaTeX input with generic forms denoting their usage, such as $\langle \text{var1} \rangle$ up to $\langle \text{varN} \rangle$ for variables,

Theorem. Consider the following series of commands such that

```
S := 3;
S := 3 + S * Z;
S := 1 + S * Z
```

Allow $Z = y$, for any natural number y , ahead of running this code then $S = 3 \times y^2 + 3 \times y + 1$ after the set of instructions has executed.

Proof. By application of usual Hoare logic:

$$\begin{array}{l} \{Z = y\} \\ S := 3; \\ \{Z = y \wedge S = 3\} \\ S := 3 + S * Z; \\ \{Z = y \wedge S = 3 \times y + 3\} \\ S := 1 + S * Z \\ \{Z = y \wedge S = 3 \times y^2 + 3 \times y + 1\} \end{array}$$

Hence, this program is shown to be correct. \square

```
Require Import String.
From PLF Require Import Imp.
From PLF Require Import Hoare.
Theorem poly_code_correct:
  forall y : nat,
  {{ Z = y }}
  S := 3;
  S := 3 + S * Z;
  S := 1 + S * Z
  {{ S = 3 * y ^ 2 + 3 * y + 1 }}.
Proof.
  intros.
  apply hoare_seq with
    (Q := (
      (Z = y /\ S = 3)
    ))%assertion).
  apply hoare_seq with
    (Q := (
      (Z = y /\ S = 3 * y + 3)
    ))%assertion).
  apply hoare_seq with
    (Q := (
      (Z = y /\ S = 3 * y^2 + 3 * y + 1)
    ))%assertion).
  all: eapply hoare_consequence_pre;
  try (apply hoare_asgn || assn_auto').
Qed.
```

Figure 5: Generated POLY example: [Left] the Hoare logic proof; [Right] the code correctness proof in Coq.

which effectively ensures *generalization to variable renaming* (Ferreira et al., 2022), $\langle \text{nat1} \rangle$ up to $\langle \text{natN} \rangle$ for numbers, or $\langle \text{def} \rangle$ for definitions, coupled with the use of a copying mechanism adapted from Gu et al. (2016). Note that a different generic token is introduced for each unique numerical constant or variable literal in the theorem and its proof, and the corresponding generic token is used in the Coq version. For example, considering the (LaTeX, Coq) pair in Figure 3, $\langle \text{nat1} \rangle$, $\langle \text{nat2} \rangle$, $\langle \text{nat3} \rangle$, and $\langle \text{nat4} \rangle$ would be used to replace the constants 2, 35, 7, and 5 respectively, everywhere in the LaTeX and Coq statements. Similarly, $\langle \text{var1} \rangle$, $\langle \text{var2} \rangle$, and $\langle \text{var3} \rangle$ were used to replace variable literals w , R , and Q . This is in contrast to using just two generic tokens $\langle \text{nat} \rangle$ and $\langle \text{var} \rangle$ everywhere, which would make all numbers coreferent and all variables coreferent. Preliminary experiments validated the utility of encoding these distinctions while maintaining the correct coreference in both LaTeX and Coq statements.

Overall, by using generic tokens for numbers, variables, and definitions, only a limited set of embeddings need to be trained and the model is forced to utilize contextual information in order to appro-

priately copy tokens into the Coq output. In this way, the model has the ability to generalize to unseen numbers or variable and definition names.

The original CopyNet (Gu et al., 2016) used an encoder-decoder architecture with a copying mechanism to calculate the probabilities of generating in-vocabulary tokens vs. copying tokens from the input sequence to the output. Our autoformalization task guarantees mutual exclusivity between generating (g) and copying (c) tokens, which allows using a simplified formula for calculating the probability of producing a token y_t at time step t . Letting \mathcal{V}_c denote the Coq vocabulary, X denote the input sequence of LaTeX tokens, and \mathcal{X} denote the collection of unique tokens in X , we calculate the probability of producing y_t as:

$$p(y_t) = \begin{cases} p(y_t, g) = \frac{1}{Z_t} e^{\psi_g(y_t)}, & y_t \in \mathcal{V}_c \\ p(y_t, c) = \frac{1}{Z_t} \sum_{x_j \in X: x_j = y_t} e^{\psi_c(x_j)}, & y_t \in \mathcal{X} \end{cases}$$

where $Z_t = \sum_{y_t \in \mathcal{V}_c} e^{\psi_g(y_t)} + \sum_{x_j \in X} e^{\psi_c(x_j)}$. The scoring functions are given by $\psi_g(y_t) = \mathbf{v}_{y_t}^\top \mathbf{W}_o \mathbf{s}_t$ and $\psi_c(x_j) = \tanh(\mathbf{h}_j^\top \mathbf{W}_c) \mathbf{s}_t$, where \mathbf{v}_{y_t} is a one-

hot encoding of y_t , \mathbf{h}_j is the hidden encoder state for the input token x_j , \mathbf{s}_t is the decoder state at step t , and \mathbf{W}_o and \mathbf{W}_c are learnable parameters.

3.2 Encoder-Decoder Architecture

We diverge from the standard Transformer architecture in a few crucial ways:

- Probabilities are calculated via $p(y_t)$ above.
- Absolute positional encodings are removed.
- Self-attention uses relative positional representations as in [Shaw et al. \(2018\)](#).
- Stacks of N encoder/decoder blocks have T recurrent passes.

All other aspects of the model remain unchanged from the original Transformer. We emphasize relative positional information over absolute in our model architecture. Preliminary evaluations on the EVEN-ODD dataset showed that Transformer models that use absolute positional encodings obtain 0% sequence-level accuracy on expression lengths that are not seen at training time. Removing reliance on absolute position resolves this type of systematic generalization. The use of relative positional encodings for the Transformer-based models was thus essential for achieving stronger systematic generalization, which also agrees with the findings of [Csordás et al. \(2021\)](#) on other NLP tasks.

4 Experimental Evaluations

To evaluate the performance of trained models, we ran two primary experiments: first on the collection of arithmetic examples, then on the collection of code correctness examples. All models are evaluated in terms of *sequence-level* accuracy, where an example is considered correctly processed only if the generated Coq sequence for both the theorem and its proof perfectly matches token by token the ground truth sequence. We also report *semantic-level* accuracy, for which the generated Coq theorem needs to attain a perfect sequence-level accuracy and the Coq engine verifies that the generated Coq proof truly proves the generated Coq theorem, regardless of whether it matches the ground truth version of the proof. This emphasizes that the model was able to capture the general meaning of the natural language proof by correctly translating the theorem and successfully proving it using the natural language version as a guide.

All experiments were performed on one NVIDIA RTX-A6000 GPU with 48GB of memory.

n	EVEN-ODD		COMPOSITES		POLY
	Seq	Sem	Seq	Sem	Both
2	99.6	99.8	76.7	97.6	100.0
3	99.4	99.6	64.6	94.2	100.0
4	99.4	99.4	56.1	93.9	82.1
5	99.2	99.6	54.9	94.4	99.2
6	98.8	98.8	57.1	94.3	45.1
7	99.1	99.5	58.5	93.4	96.5
8	93.8	94.0	53.5	88.3	15.7
9	98.6	98.6	53.7	93.7	98.2
10	7.0	7.0	1.2	1.6	35.6
11	0.0	0.0	0.0	0.0	93.5
12+	0.0	0.0	0.0	0.0	0.0

POWERS

Seq = 100.0	Sem = 100
-------------	-----------

Table 1: Sequence-level (Seq) and semantic-level (Sem) accuracy (%) on test examples, split by expression length, with the exception of POWERS.

4.1 Arithmetic Statements

We evaluate a Transformer model on the full data combining EVEN-ODD + COMPOSITES + POWERS and using both the theorem and its proof in each sequence. We tune a model with embedding and state sizes of 32, a feed forward width of 256, 4 encoder and decoder blocks with 4 recurrent passes, 4 attention heads, and a clipping value of 2 for self-attention. We trained this model over minibatches of size 20, optimized with Adam using $\beta_1 = 0.9$, $\beta_2 = 0.98$, $\varepsilon = 1e - 9$, and an initial learning rate of 0.001, annealed by a factor of $1/\sqrt{10}$ based on training loss plateaus with a patience of 5 epochs.

The results in Table 1 show that the model generalizes well to the intermediate lengths of $\{4, 6, 8\}$, with a small number of correctly translated examples longer than the maximum of 9 used in training. Otherwise, the model fails to generalize to longer unseen lengths, which is not surprising, given that Transformer models are known to fail dramatically at systematic generalization on longer inputs for various NLP tasks ([Csordás et al., 2021](#)), or to incur substantial decrease in accuracy for longer symbolic integration problems ([Welleck et al., 2022](#)). Switching to semantic-level evaluation leads to a significant increase in accuracy for COMPOSITES, with a more modest increase for EVEN-ODD.

4.2 Code Correctness Statements

We extend our scope to include data representing proofs of program correctness using the language

of Hoare logic. We train a separate model with the same embedding and state sizes, feed forward width, and learning rates as in Section 4.1. Depth is increased to 8 encoder and decoder blocks with 8 recurrent passes, 8 attention heads, and a clipping value of 8. The model is trained over minibatches of size 1 with Adam, with a patience of 3 epochs.

The POLY results shown in Table 1 demonstrate that the model is able to generalize to program line counts of $\{4, 6, 8, 10\}$ unseen during training with diminishing returns as the program length grows, eventually failing to generalize for lengths longer than the maximum seen in training. We observe that increasing the depth of the model significantly improved generalization. A model with identical hyperparameters to the arithmetic experiment yielded less than half the sequence-level accuracy for intermediate program lengths. Therefore, further increasing the depth of the model could push performance closer to optimal generalization to intermediate lengths at the cost of significantly more computing resources. Additionally, POLY examples are far less prone to non-fatal token swapping errors. We observe that semantic-level accuracy is identical to sequence-level, as all copying errors compromised the validity of the proof. Therefore, accuracies are shown as one column (Both).

4.3 Handwritten Examples

We also evaluate the semantic-level accuracy of the trained models on the collection of 45 human-written LaTeX theorem-proof pairs. This is done by manually verifying that the generated Coq theorem corresponds to the LaTeX version and that the subsequent proof is correct according to the Coq interpreter. The fully trained model achieved 53.3% for both EVEN-ODD and COMPOSITES, and 73.3% for POWERS.

Mistakes in almost all cases are confined to the mishandling of out-of-vocabulary tokens, such as mis-copying a variable within a definition or the omission of an assertion in the proof tied to a term. The model otherwise generated syntactically sound Coq code. Mistakes strongly correlate with examples that deviate significantly from the grammatical structure of the artificial data. Thus, pre-trained language models as evaluated by Wu et al. (2022) or pre-training new models on mathematical corpora like MATH (Hendrycks et al., 2021) may serve to alleviate the problems caused by the scarcity of aligned natural and formal mathematics data.

5 Concluding Remarks

As we have seen, it is feasible to train machine learning models to perform autoformalization over very restricted domains of math and code correctness proofs. These models show capability to systematically generalize to new expression lengths and program sizes. Moreover, these models were able to translate previously unseen hand written natural language examples, albeit with lower accuracy. We are hopeful that this approach can be applied to autoformalization of a larger segment of mathematics and code verification.

As mentioned by Szegedy (2020), "Autoformalization is not just a challenge: successful autoformalization would represent a breakthrough for general AI with significant implications in various domains." We see an especially significant impact in education, where integration of autoformalization into proof assistants for introductory mathematics and software verification courses would enable the detection of missing steps or misconceptions in students' proofs.

References

- K. Appel and W. Haken. 1977. [Every planar map is four colorable, part I: discharging](#). *Illinois Journal of Mathematics*, 21(3):429 – 490.
- K. Appel, W. Haken, and J. Koch. 1977. [Every planar map is four colorable. II: Reducibility](#). *Ill. J. Math.*, 21:491–567.
- Grzegorz Bancerek. 2006. Automatic translation in formalized mathematics. *Mechanized Mathematics and Its Applications*, 5(2):19–31.
- Yves Bertot and Pierre Castéran. 2013. [Interactive theorem proving and program development: Coq'Art: the calculus of inductive constructions](#). Springer Science & Business Media.
- Benjamin Andrew Carman. 2021. [Translating LaTeX to Coq: A Recurrent Neural Network Approach to Formalizing Natural Language Proofs](#). Ph.D. thesis, Ohio University.
- Róbert Csordás, Kazuki Irie, and Juergen Schmidhuber. 2021. The devil is in the detail: Simple tricks improve systematic generalization of transformers. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 619–634.
- Mostafa Dehghani, Stephan Gouws, Oriol Vinyals, Jakob Uszkoreit, and Łukasz Kaiser. 2018. Universal transformers. *arXiv preprint arXiv:1807.03819*.

- Deborah Ferreira, Mokanarangan Thayaparan, Marco Valentino, Julia Rozanova, and Andre Freitas. 2022. [To be or not to be an Integer? Encoding Variables for Mathematical Text](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 938–948, Dublin, Ireland. Association for Computational Linguistics.
- G. Gonthier. 2013. [Engineering mathematics: The odd order theorem proof](#). In *Conference Record of the Annual ACM Symposium on Principles of Programming Languages*, pages 1–2. Cited By 12.
- Georges Gonthier. 2008. Formal proof – the four-color theorem. *Notices of the AMS*, 55(11):1382–1393.
- Georges Gonthier, Andrea Asperti, Jeremy Avigad, Yves Bertot, Cyril Cohen, François Garillot, Stéphane Le Roux, Assia Mahboubi, Russell O’Connor, Sidi Ould Biha, Ioana Pasca, Laurence Rideau, Alexey Solovyev, Enrico Tassi, and Laurent Théry. 2013. [A machine-checked proof of the odd order theorem](#). In *Interactive Theorem Proving*, pages 163–179, Berlin, Heidelberg. Springer Berlin Heidelberg.
- Jiatao Gu, Zhengdong Lu, Hang Li, and Victor OK Li. 2016. [Incorporating copying mechanism in sequence-to-sequence learning](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1631–1640.
- Thomas Hales, Mark Adams, Gertrud Bauer, Tat Dat Dang, John Harrison, Le Truong Hoang, Cezary Kaliszyk, Victor Magron, Sean McLaughlin, Tat Thang Nguyen, and et al. 2017. [A formal proof of the Kepler conjecture](#). *Forum of Mathematics, Pi*, 5:e2.
- Thomas C. Hales. 2005. [A proof of the Kepler conjecture](#). *Ann. Math. (2)*, 162(3):1065–1185.
- Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. 2021. Measuring mathematical problem solving with the math dataset. *NeurIPS*.
- Nicholas J. Higham. 2020. *Handbook of Writing for the Mathematical Sciences*, third edition. Society for Industrial and Applied Mathematics, Philadelphia, PA, USA.
- Andrew David Irvine and Harry Deutsch. 2021. Russell’s Paradox. In Edward N. Zalta, editor, *The Stanford Encyclopedia of Philosophy*, Spring 2021 edition. Metaphysics Research Lab, Stanford University.
- Xavier Leroy. 2009. Formal verification of a realistic compiler. *Communications of the ACM*, 52(7):107–115.
- Tobias Nipkow, Lawrence C Paulson, and Markus Wenzel. 2002. *Isabelle/HOL: a proof assistant for higher-order logic*, volume 2283. Springer Science & Business Media.
- Benjamin C Pierce, Chris Casinghino, Marco Gaboardi, Michael Greenberg, Cătălin Hrițcu, Vilhelm Sjöberg, and Brent Yorgey. 2010. Software foundations. *Web-page: <http://www.cis.upenn.edu/bcpierce/sf/current/index.html>*.
- Benjamin C. Pierce, Arthur Azevedo de Amorim, Chris Casinghino, Marco Gaboardi, Michael Greenberg, Cătălin Hrițcu, Vilhelm Sjöberg, Andrew Tolmach, and Brent Yorgey. 2018. *Programming Language Foundations*. Software Foundations series, volume 2. Electronic textbook. Version 5.5. <http://www.cis.upenn.edu/~bcpierce/sf>.
- Peter Shaw, Jakob Uszkoreit, and Ashish Vaswani. 2018. [Self-attention with relative position representations](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 464–468, New Orleans, Louisiana. Association for Computational Linguistics.
- Christian Szegedy. 2020. [A promising path towards autoformalization and general artificial intelligence](#). In *Intelligent Computer Mathematics. CICM 2020. Lecture Notes in Computer Science*, volume 12236, pages 3–20. Springer.
- Thomas Tymoczko. 1979. [The four-color problem and its philosophical significance](#). *Journal of Philosophy*, 76(2):57–83.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. *Advances in Neural Information Processing Systems*, 30:5998–6008.
- Qingxiang Wang, Chad Brown, Cezary Kaliszyk, and Josef Urban. 2020. [Exploration of neural machine translation in autoformalization of mathematics in Mizar](#). In *Proceedings of the 9th ACM SIGPLAN International Conference on Certified Programs and Proofs, CPP 2020*, page 85–98, New York, NY, USA. Association for Computing Machinery.
- Qingxiang Wang, Cezary Kaliszyk, and Josef Urban. 2018. [First experiments with neural translation of informal to formal mathematics](#). In *Intelligent Computer Mathematics*, pages 255–270, Cham. Springer International Publishing.
- Sean Welleck, Peter West, Jize Cao, and Yejin Choi. 2022. [Symbolic Brittleness in Sequence Models: On Systematic Generalization in Symbolic Mathematics](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 36(8):8629–8637. Number: 8.
- Yuhuai Wu, Albert Q Jiang, Wenda Li, Markus N Rabe, Charles Staats, Mateja Jamnik, and Christian Szegedy. 2022. Autoformalization with large language models. *arXiv preprint arXiv:2205.12615*.