
Learning Continuous Semantic Representations of Symbolic Expressions

Miltiadis Allamanis¹ Pankajan Chanthirasegaran² Pushmeet Kohli³ Charles Sutton^{2,4}

Abstract

Combining abstract, symbolic reasoning with continuous neural reasoning is a grand challenge of representation learning. As a step in this direction, we propose a new architecture, called *neural equivalence networks*, for the problem of learning continuous semantic representations of algebraic and logical expressions. These networks are trained to represent semantic equivalence, even of expressions that are syntactically very different. The challenge is that semantic representations must be computed in a syntax-directed manner, because semantics is compositional, but at the same time, small changes in syntax can lead to very large changes in semantics, which can be difficult for continuous neural architectures. We perform an exhaustive evaluation on the task of checking equivalence on a highly diverse class of symbolic algebraic and boolean expression types, showing that our model significantly outperforms existing architectures.

1. Introduction

Combining abstract, symbolic reasoning with continuous neural reasoning is a grand challenge of representation learning. This is particularly important while dealing with exponentially large domains such as source code and logical expressions. Symbolic notation allows us to abstractly represent a large set of states that may be perceptually very different. Although symbolic reasoning is very powerful, it also tends to be hard. For example, problems such as the satisfiability of boolean expressions and automated formal proofs tend to be NP-hard or worse. This raises the exciting opportunity of using pattern recognition within symbolic reasoning, that is, to learn patterns from datasets of symbolic expressions that approximately represent se-

matic relationships. However, apart from some notable exceptions (Alemi et al., 2016; Loos et al., 2017; Zaremba et al., 2014), this area has received relatively little attention in machine learning. In this work, we explore the direction of learning continuous *semantic* representations of symbolic expressions. The goal is for expressions with similar semantics to have similar continuous representations, even if their syntactic representation is very different. Such representations have the potential to allow a new class of symbolic reasoning methods based on heuristics that depend on the continuous representations, for example, by guiding a search procedure in a symbolic solver based on a distance metric in the continuous space.

In this paper, we make a first essential step of addressing the problem of learning continuous semantic representations (SEMVECS) for symbolic expressions. Our aim is, given access to a training set of pairs of expressions for which semantic equivalence is known, to assign continuous vectors to symbolic expressions in such a way that semantically equivalent, but syntactically diverse expressions are assigned to identical (or highly similar) continuous vectors. This is an important but hard problem; learning composable SEMVECS of symbolic expressions requires that we learn about the semantics of symbolic elements and operators and how they map to the continuous representation space, thus encapsulating implicit knowledge about symbolic semantics and its recursive abstractive nature. As we show in our evaluation, relatively simple logical and polynomial expressions present significant challenges and their semantics cannot be sufficiently represented by existing neural network architectures.

Our work is similar in spirit to the work of Zaremba et al. (2014), who focus on learning expression representations to aid the search for computationally efficient identities. They use recursive neural networks (TREENN)¹ (Socher et al., 2012) for modeling *homogenous, single-variable* polynomial expressions. While they present impressive results, we find that the TREENN model fails when applied to more complex symbolic polynomial and boolean expressions. In particular, in our experiments we find that TREENNs tend to assign similar representations to syntactically similar expressions, even when they are semantically very different. The underlying conceptual problem is how to develop a con-

Work started when M. Allamanis was at Edinburgh. This work was done while P. Kohli was at Microsoft. ¹Microsoft Research, Cambridge, UK ²University of Edinburgh, UK ³DeepMind, London, UK ⁴The Alan Turing Institute, London, UK. Correspondence to: Miltiadis Allamanis <t-mialla@microsoft.com>.

¹To avoid confusion, we use TREENN for *recursive* neural networks and RNN for *recurrent* neural networks.

tinuous representation that follows syntax but *not too much*, that respects compositionality while also representing the fact that a small syntactic change can be a large semantic one.

To tackle this problem, we propose a new architecture, called *neural equivalence networks* (EQNET). EQNETs learn how syntactic composition recursively composes SEMVECS, like a TREENN, but are also designed to model large changes in semantics as the network progresses up the syntax tree. As equivalence is transitive, we formulate an objective function for training based on equivalence classes rather than pairwise decisions. The network architecture is based on composing residual-like multi-layer networks, which allows more flexibility in modeling the semantic mapping up the syntax tree. To encourage representations within an equivalence class to be tightly clustered, we also introduce a training method that we call *subexpression autoencoding*, which uses an autoencoder to force the representation of each subexpression to be predictable and reversible from its syntactic neighbors. Experimental evaluation on a highly diverse class of symbolic algebraic and boolean expression types shows that EQNETs dramatically outperform existing architectures like TREENNs and RNNs.

To summarize, the main contributions of our work are: (a) We formulate the problem of learning continuous semantic representations (SEMVECS) of symbolic expressions and develop benchmarks for this task. (b) We present neural equivalence networks (EQNETs), a neural network architecture that learns to represent expression semantics onto a continuous semantic representation space and how to perform symbolic operations in this space. (c) We provide an extensive evaluation on boolean and polynomial expressions, showing that EQNETs perform dramatically better than state-of-the-art alternatives. Code and data are available at groups.inf.ed.ac.uk/cup/semvec.

2. Model

In this work, we are interested in learning semantic, compositional representations of mathematical expressions, which we call SEMVECS, and in learning to generate identical representations for expressions that are *semantically* equivalent, *i.e.* they belong to the same equivalence class. Equivalence is a stronger property than similarity, which has been the focus of previous work in neural network learning (Chopra et al., 2005), since equivalence is additionally a transitive relationship.

Problem Hardness. Finding the equivalence of arbitrary symbolic expressions is a NP-hard problem or worse. For example, if we focus on boolean expressions, reducing an expression to the representation of the `false` equivalence class amounts to proving its non-satisfiability — an NP-complete problem. Of course, we do *not* expect to circum-

vent an NP-complete problem with neural networks. A network for solving boolean equivalence would require an exponential number of nodes in the size of the expression if $P \neq NP$. Instead, our goal is to develop architectures that efficiently learn to solve the equivalence problems for expressions that are similar to a smaller number of expressions in a given training set. The supplementary material shows a sample of such expressions that illustrate the hardness of this problem.

Notation and Framework. To allow our representations to be compositional, we employ the general framework of recursive neural networks (TREENN) (Socher et al., 2012; 2013), in our case operating on tree structures of the syntactic parse of a formula. Given a tree T , TREENNs learn distributed representations for each node in the tree by recursively combining the representations of its subtrees using a neural network. We denote the children of a node n as $\text{ch}(n)$ which is a (possibly empty) ordered tuple of nodes. We also use $\text{par}(n)$ to refer to the parent node of n . Each node in our tree has a type, *e.g.* a terminal node could be of type “a” referring to the variable a or of type “and” referring to a node of the logical AND (\wedge) operation. We refer to the type of a node n as τ_n . In pseudocode, TREENNs retrieve the representation of a tree T rooted at node ρ , by invoking the function $\text{TREENN}(\rho)$ that returns a vector representation $\mathbf{r}_\rho \in \mathbb{R}^D$, *i.e.*, a SEMVEC. The function is defined as

```
TREENN (current node  $n$ )
  if  $n$  is not a leaf then
     $\mathbf{r}_n \leftarrow \text{COMBINE}(\text{TREENN}(c_0), \dots, \text{TREENN}(c_k), \tau_n)$ ,
    where  $(c_0, \dots, c_k) = \text{ch}(n)$ 
  else
     $\mathbf{r}_n \leftarrow \text{LOOKUPLEAFEMBEDDING}(\tau_n)$ 
  return  $\mathbf{r}_n$ 
```

The general framework of TREENN allows two points of variation, the implementation of LOOKUPLEAFEMBEDDING and COMBINE. Traditional TREENNs (Socher et al., 2013) define LOOKUPLEAFEMBEDDING as a simple lookup operation within a matrix of embeddings and COMBINE as a single-layer neural network. As discussed next, these will both prove to be serious limitations in our setting. To train these networks to learn SEMVECS, we will use a supervised objective based on a set of known equivalence relations (see Section 2.2).

2.1. Neural Equivalence Networks

Our domain requires that the network learns to abstract away syntax, assigning identical representations to expressions that may be syntactically different but semantically equivalent, and also assigning different representations to expressions that may be syntactically very similar but non-equivalent. In this work, we find that standard neural architectures do not handle well this challenge. To represent semantics from syntax, we need to learn to recursively

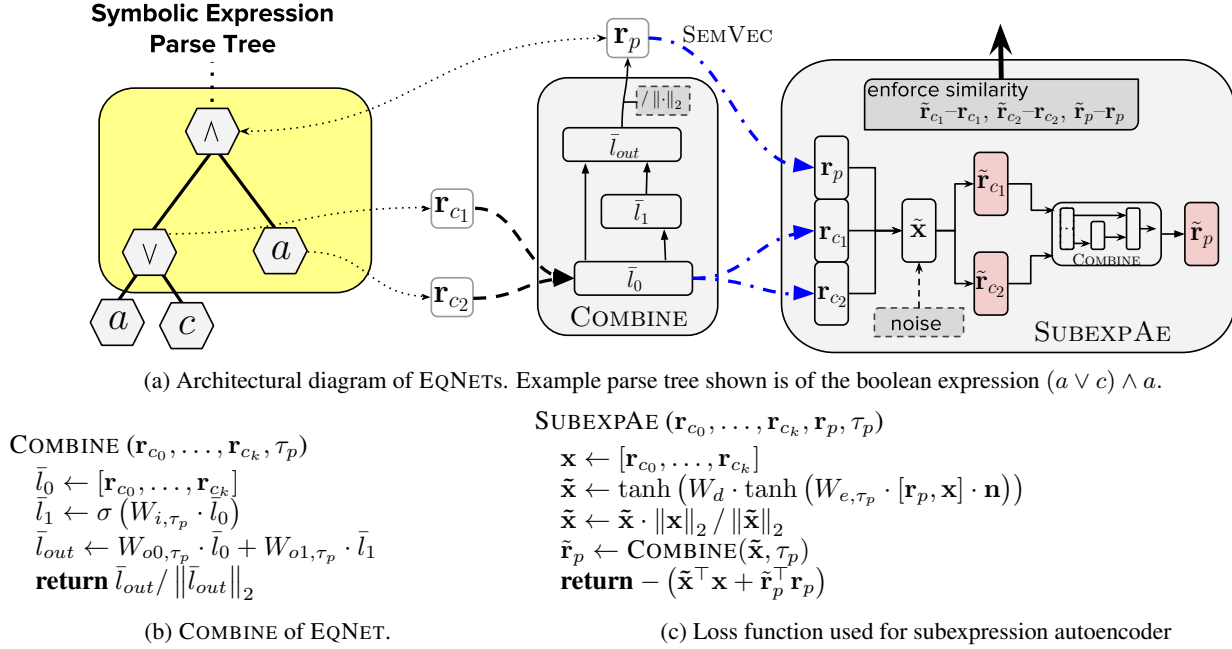


Figure 1. EQNET architecture.

compose and decompose semantic representations and remove syntactic “noise”. Any syntactic operation may significantly change semantics (*e.g.* negation, or appending \wedge FALSE) while we may reach the same semantic state through many possible operations. This necessitates using high-curvature operations over the semantic representation space. Furthermore, some operations are semantically reversible and thus we need to learn reversible semantic representations (*e.g.* $\neg\neg A$ and A should have an identical SEMVECs). Based on these, we define *neural equivalence networks* (EQNET), which learn to compose representations of equivalence classes into new equivalence classes (Figure 1a). Our network follows the TREENN architecture, *i.e.* is implemented using TREENN to model the compositional nature of symbolic expressions but is adapted based on the domain requirements. The extensions we introduce have two aims: first, to improve the network training; and second, and more interestingly, to encourage the learned representations to abstract away surface level information while retaining semantic content.

The first extension that we introduce is to the network structure at each layer in the tree. Traditional TREENNs (Socher et al., 2013) use a single-layer neural network at each tree node. During our preliminary investigations and in Section 3, we found that single layer networks are *not* adequately expressive to capture all operations that transform the input SEMVECs to the output SEMVEC and maintain semantic equivalences, requiring high-curvature operations. Part of the problem stems from the fact that within the Euclidean space of SEMVECs some operations need to be non-linear. For example a simple XOR boolean operator requires high-curvature operations in the continuous semantic representation space. Instead, we turn to multi-layer neural

networks. In particular, we define the network as shown in the function COMBINE in Figure 1b. This uses a two-layer MLP with a residual-like connection to compute the SEMVEC of each parent node in that syntax tree given that of its children. Each node type τ_n , *e.g.*, each logical operator, has a different set of weights. We experimented with deeper networks but this did not yield any improvements.

However, as TREENNs become deeper, they suffer from optimization issues, such as diminishing and exploding gradients. This is essentially because of the highly compositional nature of tree structures, where the same network (*i.e.* the COMBINE non-linear function) is used recursively, causing it to “echo” its own errors and producing unstable feedback loops. We observe this problem even with only two-layer MLPs, as the overall network can become quite deep when using two layers for each node in the syntax tree. We resolve this issue in the training procedure by constraining each SEMVEC to have unit norm. That is, we set $\text{LOOKUPLEAFEMBEDDING}(\tau_n) = C_{\tau_n} / \|C_{\tau_n}\|_2$, and we normalize the output of the final layer of COMBINE in Figure 1b. The normalization step of \bar{l}_{out} and C_{τ_n} is somewhat similar to weight normalization (Salimans & Kingma, 2016) and vaguely resembles layer normalization (Ba et al., 2016). Normalizing the SEMVECs partially resolves issues with diminishing and exploding gradients, and removes a spurious degree of freedom in the semantic representation. As simple as this modification may seem, we found it vital for obtaining good performance, and all of our multi-layer TREENNs converged to low-performing settings without it.

Although these modifications seem to improve the representation capacity of the network and its ability to be trained, we found that they were not on their own sufficient for good

performance. In our early experiments, we noticed that the networks were primarily focusing on syntax instead of semantics, *i.e.*, expressions that were nearby in the continuous space were primarily ones that were syntactically similar. At the same time, we observed that the networks did not learn to unify representations of the same equivalence class, observing multiple syntactically distinct but semantically equivalent expressions to have distant SEMVECS.

Therefore we modify the training objective in order to encourage the representations to become more abstract, reducing their dependence on surface-level syntactic information. We add a regularization term on the SEMVECS that we call a *subexpression autoencoder* (SUBEXPAE). We design this regularization to encourage the SEMVECS to have two properties: abstraction and reversibility. Because *abstraction* arguably means removing irrelevant information, a network with a bottleneck layer seems natural, but we want the training objective to encourage the bottleneck to discard syntactic information rather than semantic information. To achieve this, we introduce a component that aims to encourage *reversibility*, which we explain by an example. Observe that given the semantic representation of any two of the three nodes of a subexpression (by which we mean the parent, left child, right child of an expression tree) it is often possible to completely determine or at least place strong constraints on the semantics of the third. For example, consider a boolean formula $F(a, b) = F_1(a, b) \vee F_2(a, b)$ where F_1 and F_2 are arbitrary propositional formulae over the variables a, b . Then clearly if we know that F implies that a is true but F_1 does not, then F_2 must imply that a is true. More generally, if F belongs to some equivalence class e_0 and F_1 belongs to a different class e_1 , we want the continuous representation of F_2 to reflect that there are strong constraints on the equivalence class of F_2 .

Subexpression autoencoding encourages abstraction by employing an autoencoder with a bottleneck, thereby removing irrelevant information from the representations, and encourages reversibility by autoencoding the parent and child representations together, to encourage dependence in the representations of parents and children. More specifically, given any node p in the tree with children $c_0 \dots c_k$, we can define a parent-children tuple $[\mathbf{r}_{c_0}, \dots, \mathbf{r}_{c_k}, \mathbf{r}_p]$ containing the (computed) SEMVECS of the children and parent nodes. What SUBEXPAE does is to autoencode this representation tuple into a low-dimensional space with a denoising autoencoder. We then seek to minimize the reconstruction error of the child representations $(\tilde{\mathbf{r}}_{c_0}, \dots, \tilde{\mathbf{r}}_{c_k})$ as well as the reconstructed parent representation $\tilde{\mathbf{r}}_p$ that can be computed from the reconstructed children. More formally, we minimize the return value of SUBEXPAE in Figure 1c where \mathbf{n} is a binary noise vector with κ percent of its elements set to zero. Note that the encoder is specific to the parent node type τ_p . Although our SUBEXPAE may seem similar to the recursive autoencoders of Socher et al. (2011), it differs

in two major ways. First, SUBEXPAE autoencodes on the entire parent-children representation tuple, rather than the child representations alone. Second, the encoding is *not* used to compute the parent representation, but only serves as a regularizer.

Subexpression autoencoding has several desirable effects. First, it forces each parent-children tuple to lie in a low-dimensional space, requiring the network to compress information from the individual subexpressions. Second, because the denoising autoencoder is reconstructing parent and child representations together, this encourages child representations to be predictable from parents and siblings. Putting these two together, the goal is that the information discarded by the autoencoder bottleneck will be more syntactic than semantic, assuming that the semantics of child node is more predictable from its parent and sibling than its syntactic realization. The goal is to nudge the network to learn consistent, reversible semantics. Additionally, subexpression autoencoding has the potential to gradually unify distant representations that belong to the same equivalence class. To illustrate this point, imagine two semantically equivalent c'_0 and c''_0 child nodes of different expressions that have distant SEMVECS, *i.e.* $\|\mathbf{r}_{c'_0} - \mathbf{r}_{c''_0}\|_2 \gg \epsilon$ although $\text{COMBINE}(\mathbf{r}_{c'_0}, \dots) \approx \text{COMBINE}(\mathbf{r}_{c''_0}, \dots)$. In some cases due to the autoencoder noise, the differences between the input tuple \mathbf{x}' , \mathbf{x}'' that contain $\mathbf{r}_{c'_0}$ and $\mathbf{r}_{c''_0}$ will be non-existent and the decoder will predict a single location $\tilde{\mathbf{r}}_{c_0}$ (possibly different from $\mathbf{r}_{c'_0}$ and $\mathbf{r}_{c''_0}$). Then, when minimizing the reconstruction error, both $\mathbf{r}_{c'_0}$ and $\mathbf{r}_{c''_0}$ will be attracted to $\tilde{\mathbf{r}}_{c_0}$ and eventually should merge.

2.2. Training

We train EQNETS from a dataset of expressions whose semantic equivalence is known. Given a training set $\mathcal{T} = \{T_1 \dots T_N\}$ of parse trees of expressions, we assume that the training set is partitioned into equivalence classes $\mathcal{E} = \{e_1 \dots e_J\}$. We use a supervised objective similar to classification; the difference between classification and our setting is that whereas standard classification problems consider a fixed set of class labels, in our setting the number of equivalence classes in the training set will vary with N . Given an expression tree T that belongs to the equivalence class $e_i \in \mathcal{E}$, we compute the probability

$$P(e_i|T) = \frac{\exp(\text{TREENN}(T)^\top \mathbf{q}_{e_i} + b_i)}{\sum_j \exp(\text{TREENN}(T)^\top \mathbf{q}_{e_j} + b_j)} \quad (1)$$

where \mathbf{q}_{e_i} are model parameters that we can interpret as representations of each equivalence class that appears in the training class, and b_i are scalar bias terms. Note that in this work, we only use information about the equivalence class of the whole expression T , ignoring available information about subexpressions. This is without loss of generality, because if we do know the equivalence class of a subexpression of T , we can simply add that subexpression to

the training set. To train the model, we use a max-margin objective that maximizes classification accuracy, *i.e.*

$$\mathcal{L}_{\text{Acc}}(T, e_i) = \max \left(0, \arg \max_{e_j \neq e_i, e_j \in \mathcal{E}} \log \frac{P(e_j|T)}{P(e_i|T)} + m \right) \quad (2)$$

where $m > 0$ is a scalar margin. And therefore the optimized loss function for a single expression tree T that belongs to equivalence class $e_i \in \mathcal{E}$ is

$$\mathcal{L}(T, e_i) = \mathcal{L}_{\text{Acc}}(T, e_i) + \frac{\mu}{|Q|} \sum_{n \in Q} \text{SUBEXPAE}(\text{ch}(n), n) \quad (3)$$

where $Q = \{n \in T : |\text{ch}(n)| > 0\}$, *i.e.* contains the non-leaf nodes of T and $\mu \in (0, 1]$ a scalar weight. We found that subexpression autoencoding is counterproductive early in training, before the SEMVECS begin to represent aspects of semantics. So, for each epoch t , we set $\mu = 1 - 10^{-\nu t}$ with $\nu \geq 0$. Instead of the supervised objective that we propose, an alternative option for training EQNET would be a Siamese objective (Chopra et al., 2005) that learns about similarities (rather than equivalence) between expressions. In practice, we found the optimization to be very unstable, yielding suboptimal performance. We believe that this has to do with the compositional and recursive nature of the task that creates unstable dynamics and the fact that equivalence is a stronger property than similarity.

3. Evaluation

Datasets. We generate datasets of expressions grouped into equivalence classes from two domains. The datasets from the **BOOL** domain contain boolean expressions and the **POLY** datasets contain polynomial expressions. In both domains, an expression is either a variable, a binary operator that combines two expressions, or a unary operator applied to a single expression. When defining equivalence, we interpret distinct variables as referring to different entities in the domain, so that, e.g., the polynomials $c \cdot (a \cdot a + b)$ and $f \cdot (d \cdot d + e)$ are not equivalent. For each domain, we generate “simple” datasets which use a smaller set of possible operators and “standard” datasets which use a larger set of more complex operators. We generate each dataset by exhaustively generating *all* parse trees up to a maximum tree size. All expressions are symbolically simplified into a canonical form in order to determine their equivalence class and are grouped accordingly. Table 1 shows the datasets we generated. In the supplementary material we present some sample expressions. For the polynomial domain, we also generated **ONEV-POLY** datasets, which are polynomials over a single variable, since they are similar to the setting considered by Zaremba et al. (2014) — although **ONEV-POLY** is still a little more general because it is not restricted to homogeneous polynomials. Learning SEMVECS for boolean expressions

is already a hard problem; with n boolean variables, there are 2^{2^n} equivalence classes (*i.e.* one for each possible truth table). We split the datasets into training, validation and test sets. We create two test sets, one to measure generalization performance on equivalence classes that were seen in the training data (**SEENEQCLASS**), and one to measure generalization to unseen equivalence classes (**UNSEENEQCLASS**). It is easiest to describe **UNSEENEQCLASS** first. To create the **UNSEENEQCLASS**, we randomly select 20% of all the equivalence classes, and place all of their expressions in the test set. We select equivalence classes only if they contain at least two expressions but less than three times the average number of expressions per equivalence class. We thus avoid selecting very common (and hence trivial to learn) equivalence classes in the testset. Then, to create **SEENEQCLASS**, we take the remaining 80% of the equivalence classes, and randomly split the expressions in each class into training, validation, **SEENEQCLASS** test in the proportions 60%–15%–25%. We provide the datasets online at groups.inf.ed.ac.uk/cup/semvec.

Baselines. To compare the performance of our model, we train the following baselines. **TF-IDF**: learns a representation given the expression tokens (variables, operators and parentheses). This captures topical/declarative knowledge but is unable to capture procedural knowledge. **GRU** refers to the token-level gated recurrent unit encoder of Bahdanau et al. (2015) that encodes the token-sequence of an expression into a distributed representation. **Stack-augmented RNN** refers to the work of Joulin & Mikolov (2015) which was used to learn algorithmic patterns and uses a stack as a memory and operates on the expression tokens. We also include two recursive neural networks (**TREENN**). The **1-layer TREENN** which is the original **TREENN** also used by Zaremba et al. (2014). We also include a **2-layer TREENN**, where **COMBINE** is a classic two-layer MLP without residual connections. This shows the effect of SEMVEC normalization and subexpression autoencoder.

Hyperparameters. We tune the hyperparameters of all models using Bayesian optimization (Snoek et al., 2012) on a boolean dataset with 5 variables and maximum tree size of 7 (not shown in Table 1) using the average k -NN ($k = 1, \dots, 15$) statistics (described next). The selected hyperparameters are detailed in the supplementary material.

3.1. Quantitative Evaluation

Metrics. To evaluate the quality of the learned representations we count the proportion of k nearest neighbors of each expression (using cosine similarity) that belong to the same equivalence class. More formally, given a test query expression q in an equivalence class c we find the k nearest neighbors $\mathbb{N}_k(q)$ of q across all expressions, and define the

Table 1. Dataset statistics and results. SIMP datasets contain simple operators (“ \wedge , \vee , \neg ” for BOOL and “ $+$, $-$ ” for POLY) while the rest contain all operators (i.e. “ \wedge , \vee , \neg , \oplus , \Rightarrow ” for BOOL and “ $+$, $-$, \cdot ” for POLY). \oplus is the XOR operator. The number in the dataset name indicates its expressions’ maximum tree size. L refers to a “larger” number of 10 variables. H is the entropy of equivalence classes.

Dataset	# Vars	# Equiv Classes	# Exprs	H	tf-idf	GRU	$score_5$ (%) in UNSEENEQCLASS			
							StackRNN	1L TREENN	2L TREENN	EQNET
SIMPBOOL8	3	120	39,048	5.6	17.4	30.9	26.7	27.4	25.5	97.4
SIMPBOOL10 ^S	3	191	26,304	7.2	6.2	11.0	7.6	25.0	93.4	99.1
BOOL5	3	95	1,239	5.6	34.9	35.8	12.4	16.4	26.0	65.8
BOOL8	3	232	257,784	6.2	10.7	17.2	16.0	15.7	15.4	58.1
BOOL10 ^S	10	256	51,299	8.0	5.0	5.1	3.9	10.8	20.2	71.4
SIMPBOOLL5	10	1,342	10,050	9.9	53.1	40.2	50.5	3.48	19.9	85.0
BOOLL5	10	7,312	36,050	11.8	31.1	20.7	11.5	0.1	0.5	75.2
SIMPPLY5	3	47	237	5.0	21.9	6.3	1.0	40.6	27.1	65.6
SIMPPLY8	3	104	3,477	5.8	36.1	14.6	5.8	12.5	13.1	98.9
SIMPPLY10	3	195	57,909	6.3	25.9	11.0	6.6	19.9	7.1	99.3
ONEV-POLY10	1	83	1,291	5.4	43.5	10.9	5.3	10.9	8.5	81.3
ONEV-POLY13	1	677	107,725	7.1	3.2	4.7	2.2	10.0	56.2	90.4
POLY5	3	150	516	6.7	37.8	34.1	2.2	46.8	59.1	55.3
POLY8	3	1,102	11,451	9.0	13.9	5.7	2.4	10.4	14.8	86.2

^S dataset contains all equivalence classes but at most 200 uniformly sampled (without replacement) expressions per equivalence class.

score as

$$score_k(q) = \frac{|\mathbb{N}_k(q) \cap c|}{\min(k, |c|)}. \quad (4)$$

To report results for a given testset, we simply average $score_k(q)$ for all expressions q in the testset. We also report the precision-recall curves for the problem of clustering the SEMVECS into their appropriate equivalence classes.

Evaluation. Figure 2 presents the average per-model precision-recall curves across the datasets. Table 1 shows $score_5$ of UNSEENEQCLASS. Detailed plots are found in the supplementary material. EQNET performs better for all datasets, by a large margin. The only exception is POLY5, where the 2-L TREENN performs better. However, this may have to do with the small size of the dataset. The reader may observe that the simple datasets (containing fewer operations and variables) are easier to learn. Understandably, introducing more variables increases the size of the represented space reducing performance. The tf-idf method performs better in settings with more variables, because it captures well the variables and operations used. Similar observations can be made for sequence models. The one and two layer TREENN have mixed performance; we believe that this has to do with exploding and diminishing gradients due to the deep and highly compositional nature of TREENN. Although Zaremba et al. (2014) consider a different problem to us, they use data similar to the ONEV-POLY datasets with a traditional TREENN architecture. Our evaluation suggests that EQNETS perform much better within the ONEV-POLY setting.

Evaluation of Compositionality. We evaluate whether EQNETS successfully learn to compute compositional representations, rather than overfitting to expression trees of

a small size. To do this we consider a type of transfer setting, in which we train on simpler datasets, but test on more complex ones; for example, training on the training set of BOOL5 but testing on the testset of BOOL8. We average over 11 different train-test pairs (full list in supplementary material) and show the results in Figure 3a and Figure 3b. These graphs again show that EQNETS are better than any of the other methods, and indeed, performance is only a bit worse than in the non-transfer setting.

Impact of EQNET Components EQNETS differ from traditional TREENN in two major ways, which we analyze here. First, SUBEXP AE improves performance. When training the network with and without SUBEXP AE, on average, the area under the curve (AUC) of $score_k$ decreases by 16.8% on the SEENEQCLASS and 19.7% on the UNSEENEQCLASS. This difference is smaller in the transfer setting, where AUC decreases by 8.8% on average. However, even in this setting we observe that SUBEXP AE helps more in large and diverse datasets. The second key difference to traditional TREENN is the output normalization and the residual connections. Comparing our model to the one-layer and two-layer TREENN again, we find that output normalization results in important improvements (the two-layer TREENN have on average 60.9% smaller AUC). We note that only the combination of the residual connections and the output normalization improve the performance, whereas when used separately, there are no significant improvements over the two-layer TREENN.

3.2. Qualitative Evaluation

Table 2 shows expressions whose SEMVEC nearest neighbor is of an expression of another equivalence class. Manually inspecting boolean expressions, we find that EQNET confusions happen more when a XOR or implication operator is


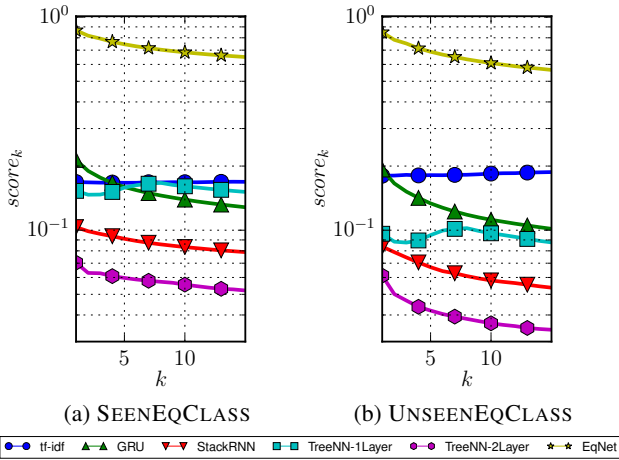


Figure 2. Precision-Recall Curves averaged across datasets.


Table 2. Non semantically equivalent first nearest-neighbors from BOOL8 and POLY8. A checkmark indicates that the method correctly results in the nearest neighbor being from the same equivalence class.

Expr	$a \wedge (a \wedge (a \wedge (\neg c)))$	$a \wedge (a \wedge (c \Rightarrow (\neg c)))$	$(a \wedge a) \wedge (c \Rightarrow (\neg c))$	$a + (c \cdot (a + c))$	$((a + c) \cdot c) + a$	$(b \cdot b) - b$
tfidf	$c \wedge ((a \wedge a) \wedge (\neg a))$	$c \Rightarrow (\neg((c \wedge a) \wedge a))$	$c \Rightarrow (\neg((c \wedge a) \wedge a))$	$a + (c + a) \cdot c$	$(c \cdot a) + (a + c)$	$b \cdot (b - b)$
GRU	✓	$a \wedge (a \wedge (c \wedge (\neg c)))$	$(a \wedge a) \wedge (c \Rightarrow (\neg c))$	$b + (c \cdot (a + c))$	$((b + c) \cdot c) + a$	$(b + b) \cdot b - b$
1L-TREENN	$a \wedge (a \wedge (a \wedge (\neg b)))$	$a \wedge (a \wedge (c \Rightarrow (\neg b)))$	$(a \wedge a) \wedge (c \Rightarrow (\neg b))$	$a + (c \cdot (b + c))$	$((b + c) \cdot c) + a$	$(a - c) \cdot b - b$
EQNET	✓	✓	$(\neg(b \Rightarrow (b \vee c))) \wedge a$	✓	✓	$(b \cdot b) \cdot b - b$


 Figure 3. Evaluation of compositionality; training set simpler than test set. Average $score_k$ (y -axis in log-scale). Markers are shown every three ticks for clarity. TREENN refers to Socher et al. (2012).

involved. In fact, we fail to find any confused expressions for EQNET *not* involving these operations in BOOL5 and in the top 100 expressions in BOOL10. As expected, tf-idf confuses expressions with others that contain the same operators and variables ignoring order. In contrast, GRU and TREENN tend to confuse expressions with very similar symbolic representations, *i.e.* that differ in one or two deeply nested variables or operators. In contrast, EQNET tends to confuse fewer expressions (as we previously showed) and the confused expressions tend to be more syntactically diverse and semantically related.

Figure 4 shows a visualization of $score_5$ for each node in the expression tree. One may see that as EQNET knows how


 Figure 4. Visualization of $score_5$ for all expression nodes for three BOOL10 and four POLY8 test sample expressions using EQNET. The darker the color, the lower the score, *i.e.* white implies a score of 1 and dark red a score of 0.

to compose expressions that achieve good score, even if the subexpressions achieve a worse score. This suggests that for common expressions, (*e.g.* single variables and monomials) the network tends to select a unique location, without merging the equivalence classes or affecting the upstream performance of the network. Larger scale interactive t-SNE visualizations can be found [online](#).

Figure 5 presents two PCA visualizations of the SEMVECS of simple expressions and their negations/negatives. It can be discerned that the black dots and their negations (in red) are discriminated in the semantic representation space. Figure 5b shows this property in a clear manner: left-right discriminates between polynomials with 1 and $-a$, top-bottom between polynomials with $-b$ and b and the diagonal parallel to $y = -x$ between c and $-c$. We observe a similar behavior in Figure 5a for boolean expressions.

4. Related Work

Researchers have proposed compilation schemes that can transform any given program or expression to an equivalent neural network (Gruau et al., 1995; Neto et al., 2003; Siegel-

Acknowledgments

This work was supported by Microsoft Research through its PhD Scholarship Programme and the Engineering and Physical Sciences Research Council [grant number EP/K024043/1]. We thank the University of Edinburgh Data Science EPSRC Centre for Doctoral Training for providing additional computational resources.

References

- Alemi, Alex A, Chollet, Francois, Irving, Geoffrey, Szegedy, Christian, and Urban, Josef. DeepMath – Deep sequence models for premise selection. *arXiv preprint arXiv:1606.04442*, 2016.
- Allamanis, Miltiadis, Peng, Hao, and Sutton, Charles. A convolutional attention network for extreme summarization of source code. In *ICML*, 2016.
- Ba, Jimmy Lei, Kiros, Jamie Ryan, and Hinton, Geoffrey E. Layer normalization. *arXiv preprint arXiv:1607.06450*, 2016.
- Bahdanau, Dzmitry, Cho, Kyunghyun, and Bengio, Yoshua. Neural machine translation by jointly learning to align and translate. In *ICLR*, 2015.
- Chopra, Sumit, Hadsell, Raia, and LeCun, Yann. Learning a similarity metric discriminatively, with application to face verification. In *CVPR*, 2005.
- Dyer, Chris, Ballesteros, Miguel, Ling, Wang, Matthews, Austin, and Smith, Noah A. Transition-based dependency parsing with stack long short-term memory. In *ACL*, 2015.
- Graves, Alex, Wayne, Greg, and Danihelka, Ivo. Neural Turing machines. *arXiv preprint arXiv:1410.5401*, 2014.
- Grefenstette, Edward, Hermann, Karl Moritz, Suleyman, Mustafa, and Blunsom, Phil. Learning to transduce with unbounded memory. In *NIPS*, 2015.
- Gruau, Frédéric, Ratajszczak, Jean-Yves, and Wiber, Gilles. A neural compiler. *Theoretical Computer Science*, 1995.
- Joulin, Armand and Mikolov, Tomas. Inferring algorithmic patterns with stack-augmented recurrent nets. In *NIPS*, 2015.
- Kaiser, Łukasz and Sutskever, Ilya. Neural GPUs learn algorithms. In *ICLR*, 2016.
- Loos, Sarah, Irving, Geoffrey, Szegedy, Christian, and Kaliszky, Cezary. Deep network guided proof search. *arXiv preprint arXiv:1701.06972*, 2017.
- Mou, Lili, Li, Ge, Zhang, Lu, Wang, Tao, and Jin, Zhi. Convolutional neural networks over tree structures for programming language processing. In *AAAI*, 2016.
- Neelakantan, Arvind, Le, Quoc V, and Sutskever, Ilya. Neural programmer: Inducing latent programs with gradient descent. In *ICLR*, 2015.
- Neto, João Pedro, Siegelmann, Hava T, and Costa, J Félix. Symbolic processing in neural networks. *Journal of the Brazilian Computer Society*, 8(3):58–70, 2003.
- Piech, Chris, Huang, Jonathan, Nguyen, Andy, Phulsuksombati, Mike, Sahami, Mehran, and Guibas, Leonidas J. Learning program embeddings to propagate feedback on student code. In *ICML*, 2015.
- Reed, Scott and de Freitas, Nando. Neural programmer-interpreters. *ICLR*, 2016.
- Salimans, Tim and Kingma, Diederik P. Weight normalization: A simple reparameterization to accelerate training of deep neural networks. In *Advances in Neural Information Processing Systems*, 2016.
- Siegelmann, Hava T. Neural programming language. In *Proceedings of the 12th National Conference on Artificial Intelligence*, 1994.
- Snoek, Jasper, Larochelle, Hugo, and Adams, Ryan P. Practical Bayesian optimization of machine learning algorithms. In *NIPS*, 2012.
- Socher, Richard, Pennington, Jeffrey, Huang, Eric H, Ng, Andrew Y, and Manning, Christopher D. Semi-supervised recursive autoencoders for predicting sentiment distributions. In *EMNLP*, 2011.
- Socher, Richard, Huval, Brody, Manning, Christopher D, and Ng, Andrew Y. Semantic compositionality through recursive matrix-vector spaces. In *EMNLP*, 2012.
- Socher, Richard, Perelygin, Alex, Wu, Jean Y, Chuang, Jason, Manning, Christopher D, Ng, Andrew Y, and Potts, Christopher. Recursive deep models for semantic compositionality over a sentiment treebank. In *EMNLP*, 2013.
- Zaremba, Wojciech, Kurach, Karol, and Fergus, Rob. Learning to discover efficient mathematical identities. In *Advances in Neural Information Processing Systems*, pp. 1278–1286, 2014.