
Mapping Natural-language Problems to Formal-language Solutions Using Structured Neural Representations

Kezhen Chen^{1,2} Qiuyuan Huang¹ Hamid Palangi¹ Paul Smolensky^{1,3} Kenneth D. Forbus² Jianfeng Gao¹

Abstract

Generating formal-language programs represented by relational tuples, such as Lisp programs or mathematical operations, to solve problems stated in natural language is a challenging task because it requires explicitly capturing discrete symbolic structural information implicit in the input. However, most general neural sequence models do not explicitly capture such structural information, limiting their performance on these tasks. In this paper, we propose a new encoder-decoder model based on a structured neural representation, *Tensor Product Representations (TPRs)*, for mapping Natural-language problems to Formal-language solutions, called **TP-N2F**. The encoder of TP-N2F employs TPR ‘binding’ to encode natural-language symbolic structure in vector space and the decoder uses TPR ‘unbinding’ to generate, in symbolic space, a sequential program represented by relational tuples, each consisting of a relation (or operation) and a number of arguments. TP-N2F considerably outperforms LSTM-based seq2seq models on two benchmarks and creates new state-of-the-art results. Ablation studies show that improvements can be attributed to the use of structured TPRs explicitly in both the encoder and decoder. Analysis of the learned structures shows how TPRs enhance the interpretability of TP-N2F.

1. Introduction

When people perform explicit reasoning, they can typically describe the way to the conclusion step by step via relational

¹Microsoft Research, Redmond, USA. ²Department of Computer Science, Northwestern University, Evanston, USA. ³Department of Cognitive Science, Johns Hopkins University, Baltimore, USA.. Correspondence to: Kezhen Chen <kzchen@u.northwestern.edu>.

descriptions. There is ample evidence that relational, structured representations are important for human cognition, e.g., (Goldin-Meadow & Gentner, 2003; Forbus et al., 2017; Crouse et al., 2018; Chen & Forbus, 2018; Chen et al., 2019; Lee et al., 2019). Although a rapidly growing number of researchers use deep learning to solve complex symbolic reasoning and language tasks (a recent review is Gao et al. (2019)), most existing deep learning models, including sequence models such as LSTMs, do not explicitly capture human-like relational structured information.

In this paper we propose a novel neural architecture, **TP-N2F**, for mapping a Natural-language (NL) question to a Formal-language (FL) program represented by a sequence of relational tuples (N2F). In the tasks we study, math or programming problems are stated in natural language, and answers are given as programs: sequences of relational structured representations, to solve the problems step by step like a human being, instead of directly generating the final answer. For example, from one of our datasets, MathQA: given a natural-language math problem “20 is subtracted from 60 percent of a number, the result is 88. Find the number?”, the formal-language solution program is “(add,n0,n2) (divide,n1,const100) (divide,#0,#1)”, where n1 indicates the first number mentioned in the question and #i indicates the output of the i^{th} previous tuple. TP-N2F encodes the natural-language symbolic structure of the problem in an input vector space, maps this to a vector in an intermediate space, and uses that vector to produce a sequence of output vectors that are decoded as relational structures. Both input and output structures are modeled as Tensor Product Representations (TPRs) (Smolensky, 1990) and the structured representations of inputs are mapped to the structured representations of outputs. During encoding, NL-input symbolic structures are encoded as vector space embeddings using TPR ‘binding’ (following Palangi et al. (2018)); during decoding, symbolic constituents are extracted from structure-embedding output vectors using TPR ‘unbinding’ (following Huang et al. (2018; 2019)). By employing TPRs, the model achieves better performance and increased interpretability.

Our contributions in this work are as follows. (i) We introduce the notion of abstract role-level analysis, and propose such an analysis of N2F tasks. (ii) We present a new TP-N2F

model which gives a neural-network-level implementation of a model solving the N2F task under the role-level description proposed in (i). To our knowledge, this is the first model to be proposed which combines both the binding and unbinding operations of TPRs to solve generation tasks through deep learning. (iii) State-of-the-art performance on two recently developed N2F tasks shows that the TP-N2F model has significant structure learning ability on tasks requiring symbolic reasoning through program synthesis.

2. Related Work

N2F tasks include many different subtasks such as symbolic reasoning or semantic parsing (Kamath & Das, 2019; Cai & Lam, 2019; Liao et al., 2018; Amini et al., 2019; Polosukhin & Skidanov, 2018; Bednarek et al., 2019). These tasks require models with strong structure-learning ability. TPR is a promising technique for encoding symbolic structural information and modeling symbolic reasoning in vector space. TPR binding has been used for encoding and exploring grammatical structural information of natural language (Palangi et al., 2018; Huang et al., 2019). TPR unbinding has also been used to generate natural language captions from images (Huang et al., 2018). Some researchers use TPRs for modeling deductive reasoning processes both on a rule-based model and deep learning models in vector space (Lee et al., 2016; Smolensky et al., 2016; Schlag & Schmidhuber, 2018). However, none of these previous models takes advantage of combining TPR binding and TPR unbinding to learn structure representation mappings explicitly, as done in our model. Although researchers are paying increasing attention to N2F tasks, most of the proposed models either do not encode structural information explicitly or are specialized to particular tasks. Our proposed TP-N2F neural model is general and can be applied to many tasks.

TP-N2F represents inputs and outputs as structures and learns to map these structures. In cognitive science and psychology, mapping one domain to another is also an important field. For example, Goldin-Meadow & Gentner (2003) proposed the Structure Mapping Theory to model human analogy within cognitive science, and Forbus et al. (2017) introduced the computational implementation, the Structure Mapping Engine (SME), of the Structure Mapping Theory. Following these works, Crouse et al. (2018); Chen & Forbus (2018); Chen et al. (2019) applied SME to language and vision problems. Researchers also explore the use of concept theory to map structural representations from different domains (Roads & Love, 2019; Martin, 2020). In this paper, we propose to use the structure-to-structure scheme to build neural models: the TP-N2F model follows this scheme.

3. Structured Representations using TPRs

The Tensor Product Representation (TPR) mechanism is a method to create a vector space embedding of complex symbolic structures. The type of a symbol structure is defined by a set of structural positions or roles, such as the left-child-of-root position in a tree, or the second-argument-of- R position of a given relation R . In a particular instance of a structural type, each of these roles may be occupied by a particular filler, which can be an atomic symbol or a substructure (e.g., the entire left sub-tree of a binary tree can serve as the filler of the role left-child-of-root). For now, we assume the fillers to be atomic symbols.¹

The TPR embedding of a symbol structure is the sum of the embeddings of all its constituents, each constituent comprising a role together with its filler. The embedding of a constituent is constructed from the embedding of a role and the embedding of the filler of that role: these are joined together by the TPR ‘binding’ operation, the tensor (or generalized outer) product \otimes .

Formally, suppose a symbolic type is defined by the roles $\{r_i\}$, and suppose that in a particular instance of that type, S , role r_i is bound by filler f_i . The TPR embedding of S is the order-2 tensor

$$\mathbf{T} = \sum_i \mathbf{f}_i \otimes \mathbf{r}_i = \sum_i \mathbf{f}_i \mathbf{r}_i^\top \quad (1)$$

where $\{\mathbf{f}_i\}$ are vector embeddings of the fillers and $\{\mathbf{r}_i\}$ are vector embeddings of the roles. In Eq. 1, and below, for notational simplicity we conflate order-2 tensors and matrices.

A TPR scheme for embedding a set of symbol structures is defined by a decomposition of those structures into roles bound to fillers, an embedding of each role as a **role vector**, and an embedding of each filler as a **filler vector**. Let the total number of roles and fillers available be n_R, n_F , respectively. Define the matrix of all possible role vectors to be $\mathbf{R} \in \mathbb{R}^{d_R \times n_R}$, with column i , $[\mathbf{R}]_{:i} = \mathbf{r}_i \in \mathbb{R}^{d_R}$, comprising the embedding of r_i . Similarly let $\mathbf{F} \in \mathbb{R}^{d_F \times n_F}$ be the matrix of all possible filler vectors. The TPR $\mathbf{T} \in \mathbb{R}^{d_F \times d_R}$. Below, d_R, n_R, d_F, n_F will be hyper-parameters, while \mathbf{R}, \mathbf{F} will be learned parameter matrices.

Using summation in Eq.1 to combine the vectors embedding the constituents of a structure risks non-recoverability of those constituents given the embedding \mathbf{T} of the structure as a whole. The tensor product is chosen as the binding operation in order to enable recovery of the filler of any role in a structure S given its TPR \mathbf{T} . This can be done with

¹When fillers are structures themselves, binding can be used recursively, giving tensors of order higher than 2. In general, binding is done with the tensor product, since conflation with matrix algebra is only possible for order-2 tensors. Our unbinding of relational tuples involves the order-3 TPRs defined in Sec. 4.1.

perfect precision if the embeddings of the roles are linearly independent. In that case the role matrix \mathbf{R} has a left inverse \mathbf{U} : $\mathbf{UR} = \mathbf{I}$. Now define the **unbinding** (or **dual**) **vector** for role r_j , \mathbf{u}_j , to be the j^{th} column of \mathbf{U}^\top : $\mathbf{U}_{:,j}^\top$. Then, since $[\mathbf{I}]_{ji} = [\mathbf{UR}]_{ji} = \mathbf{U}_{j,:} \mathbf{R}_{:,i} = [\mathbf{U}_{:,j}^\top]^\top \mathbf{R}_{:,i} = \mathbf{u}_j^\top \mathbf{r}_i = \mathbf{r}_i^\top \mathbf{u}_j$, we have $\mathbf{r}_i^\top \mathbf{u}_j = \delta_{ji}$. This means that, to recover the filler of r_j in the structure with TPR \mathbf{T} , we can take its tensor inner product (or matrix-vector product) with \mathbf{u}_j :²

$$\mathbf{T}\mathbf{u}_j = \left[\sum_i \mathbf{f}_i \mathbf{r}_i^\top \right] \mathbf{u}_j = \sum_i \mathbf{f}_i \delta_{ij} = \mathbf{f}_j \quad (2)$$

In the architecture proposed here, we make use of TPR ‘binding’ for the structured embedding encoding the natural-language problem statement; we use TPR ‘unbinding’ of the structured output embedding to decode the formal-language solution programs represented by relational tuples. Because natural-language and formal-language pertain to different representations (natural-language is an order-2 tensor and formal-language is an order-3 tensor), the NL-binding and FL-unbinding vectors are not related to one another. The structured neural Tensor Product Representations of natural-language and formal-language, and the details of binding and unbinding process used in the architecture, will be introduced in 4.1.

4. TP-N2F Model

We propose a general TP-N2F neural network architecture operating over TPRs to solve N2F tasks under a proposed role-level description of those tasks. In this description, natural-language input is represented as a straightforward order-2 tensor role structure, and formal-language relational representations of outputs are represented with a new order-3 tensor recursive role structure proposed here. Figure 1 shows an overview diagram of the TP-N2F model. It depicts the following high-level description.

As shown in Figure 1, while the natural-language input is a sequence of words, the output is a sequence of multi-argument relational tuples such as $(R A_1 A_2)$, a 3-tuple consisting of a binary relation (or operation) R with its two arguments. The “TP-N2F encoder” uses two LSTMs to produce a pair consisting of a filler vector and a role vector, which are bound together with the tensor product. These tensor products, concatenated, comprise the “context” over which attention will operate in the decoder. The sum of the word-level TPRs, flattened to a vector, is treated as a representation of the entire problem statement; it is fed to

²When the role vectors are not linearly independent, this operation performs unbinding approximately, taking \mathbf{U} to be the left pseudo-inverse of \mathbf{R} . Because randomly chosen vectors on the unit sphere in a high-dimensional space are approximately orthogonal, the approximation is often excellent.

the “Reasoning MLP”, which transforms this encoding of the problem into a vector encoding the solution. This is the initial state of the “TP-N2F decoder” attentional LSTM, which outputs at each time step an order-3 tensor representing a relational tuple. To generate a correct tuple from decoder operations, the model must learn to give the order-3 tensor the form of a TPR for a $(R A_1 A_2)$ tuple (detailed explanation in Sec. 4.1). In the following sections, we first introduce the details of our proposed role-level description for N2F tasks, and then present how our proposed TP-N2F model uses TPR binding and unbinding operations to create a neural network implementation of this description of N2F tasks.

4.1. Role-level description of N2F tasks

In this section, we propose a role-level description of N2F tasks, which specifies the filler/role structures of the input natural-language symbolic expressions and the output relational representations. As the two structures are different, we also propose a formal scheme for structure mapping on TPRs.

Role-Level Description for Natural-Language Input

Instead of encoding each token of a sentence with a non-compositional embedding vector looked up in a learned dictionary, we use a learned role-filler decomposition to compose a tensor representation for each token. Given a sentence S with n word tokens $\{w^0, w^1, \dots, w^{n-1}\}$, each word token w^t is assigned a learned role vector \mathbf{r}^t , soft-selected from the learned dictionary \mathbf{R} , and a learned filler vector \mathbf{f}^t , soft-selected from the learned dictionary \mathbf{F} (Sec. 3). The mechanism closely follows that of Palangi et al. (2018), and we hypothesize similar results: the role and filler approximately encode the structural role of the token and its lexical semantics, respectively.³ Then each word token w^t is represented by the tensor product of the role vector and the filler vector: $\mathbf{T}^t = \mathbf{f}^t \otimes \mathbf{r}^t$. In addition to the set of all its token embeddings $\{\mathbf{T}^0, \dots, \mathbf{T}^{n-1}\}$, the sentence S as a whole is assigned a TPR equal to the sum of the TPR embeddings of all its word tokens: $\mathbf{T}_S = \sum_{t=0}^{n-1} \mathbf{T}^t$.

Using TPRs to encode natural language has several advantages. First, natural language TPRs can be interpreted by exploring the distribution of tokens grouped by the role and filler vectors they are assigned by a trained model (as in Palangi et al. (2018)). Second, TPRs avoid the Bag of Word (BoW) confusion (Huang et al., 2018): the BoW encoding of *Jay saw Kay* is the same as the BoW encoding of *Kay saw Jay* but the encodings are different with TPR embedding,

³Although the TPR formalism treats fillers and roles symmetrically, in use, hyperparameters are selected so that the number of available fillers is greater than that of roles. Thus, on average, each role is assigned to more words, encouraging it to take on a more general function, such as a grammatical role.




Figure 1. Overview diagram of TP-N2F.

because the role filled by a symbol changes with its context.

Role-Level Description for Relational Representations

In this section, we propose a novel recursive role-level description for representing symbolic relational tuples. Each relational tuple contains a relation token and multiple argument tokens. Given a binary relation rel , a relational tuple can be written as $(rel \ arg_1 \ arg_2)$ where arg_1, arg_2 indicate two arguments of relation rel . Let us adopt the two positional roles, $p_i^{rel} = arg_i\text{-of-}rel$ for $i = 1, 2$. The filler of role p_i^{rel} is arg_i . Now let us use role decomposition recursively, noting that the role p_i^{rel} can itself be decomposed into a sub-role $p_i = arg_i\text{-of-}$ which has a sub-filler rel . Suppose that arg_i, rel, p_i are embedded as vectors $\mathbf{a}_i, \mathbf{r}, \mathbf{p}_i$. Then the TPR encoding of p_i^{rel} is $\mathbf{r}_{rel} \otimes \mathbf{p}_i$, so the TPR encoding of filler arg_i bound to role p_i^{rel} is $\mathbf{a}_i \otimes (\mathbf{r}_{rel} \otimes \mathbf{p}_i)$. The tensor product is associative, so we can omit parentheses and write the TPR for the formal-language expression, the relational tuple $(rel \ arg_1 \ arg_2)$, as:

$$\mathbf{H} = \mathbf{a}_1 \otimes \mathbf{r}_{rel} \otimes \mathbf{p}_1 + \mathbf{a}_2 \otimes \mathbf{r}_{rel} \otimes \mathbf{p}_2. \quad (3)$$

Given the unbinding vectors \mathbf{p}'_i for positional sub-role vectors \mathbf{p}_i and the unbinding relational vector \mathbf{r}'_{rel} for the relational vector \mathbf{r}_{rel} that embeds relation rel , each argument can be unbound in two steps as shown in Eqs. 4–5.

$$\mathbf{H} \cdot \mathbf{p}'_i = \mathbf{a}_i \otimes \mathbf{r}_{rel} \quad (4)$$

$$[\mathbf{a}_i \otimes \mathbf{r}_{rel}] \cdot \mathbf{r}'_{rel} = \mathbf{a}_i \quad (5)$$

Here \cdot denotes the tensor inner product, which for the order-3 tensor \mathbf{H} and order-1 \mathbf{p}'_i in Eq. 4 can be defined as $[\mathbf{H} \cdot \mathbf{p}'_i]_{jk} = \sum_l [\mathbf{H}]_{jkl} [\mathbf{p}'_i]_l$; in Eq. 5, \cdot is equivalent to the matrix-vector product.

Our proposed scheme can be contrasted with the TPR scheme in which $(rel \ arg_1 \ arg_2)$ is embedded as $\mathbf{r}_{rel} \otimes \mathbf{a}_1 \otimes \mathbf{a}_2$, e.g., (Smolensky et al., 2016; Schlag & Schmidhuber, 2018). In that scheme, an n -ary-relation tuple is

embedded as an order- $(n+1)$ tensor, and unbinding an argument requires knowing all the other arguments (to use their unbinding vectors). In the scheme proposed here, an n -ary-relation tuple is still embedded as an order-3 tensor: there are just n terms in the sum in Eq. 3, using n positional sub-role vectors $\mathbf{p}_1, \dots, \mathbf{p}_n$; unbinding simply requires knowing the unbinding vectors for these fixed position vectors.

In the model, the order-3 tensor \mathbf{H} of Eq. 3 has a different status than the order-2 tensor \mathbf{T}_S of Sec. 4.1. \mathbf{T}_S is a TPR by construction, whereas \mathbf{H} is a TPR as a result of successful learning. To generate the output relational tuples, the decoder assumes each tuple has the form of Eq. 3, and performs the unbinding operations which that structure calls for. In section 4.4, it is shown that, if unbinding each of a set of roles from some unknown tensor \mathbf{T} gives a target set of fillers, then \mathbf{T} must equal the TPR generated by those role/filler pairs, plus some tensor that is irrelevant because unbinding from it produces the zero vector. In other words, if the decoder succeeds in producing filler vectors that correspond to output relational tuples that match the target, then, as far as what the decoder can see, the tensor that it operates on is the TPR of Eq. 3.

TP-N2F Scheme for Learning Input-Output Mapping

To generate formal relational tuples from natural-language descriptions, a learning strategy for the mapping between the two structures is particularly important. As shown in (6), we formalize the learning scheme as learning a mapping function $f_{\text{mapping}}(\cdot)$, which, given a structural representation of the natural-language input, \mathbf{T}_S , outputs a tensor \mathbf{T}_F from which the structural representation of the output can be generated. At the role level of description, there's nothing more to be said about this mapping; how it is modeled at the neural network level is discussed in Sec. 4.2.

$$\mathbf{T}_F = f_{\text{mapping}}(\mathbf{T}_S) \quad (6)$$




Figure 2. Implementation of the TP-N2F encoder.

4.2. TP-N2F Model for N2F Generation

In this section, we introduce the TP-N2F model. The encoder and decoder are described first. Then, the inference and learning strategy are presented. Finally, we prove that successful learning entails that the tensor that is input to the decoder’s unbinding module is a TPR.

Natural-Language Encoder in TP-N2F As shown in Figure 1, the TP-N2F model is implemented with three steps: encoding, mapping, and decoding. The encoding step is implemented by the TP-N2F natural-language encoder (*TP-N2F Encoder*), which takes the sequence of word tokens as inputs, and encodes them via TPR binding according to the TP-N2F role scheme for natural-language input given in Sec. 4.1. The mapping step is implemented by an MLP called the *Reasoning Module*, which takes the encoding produced by the TP-N2F Encoder as input. It learns to map the natural-language-structure encoding of the input to a representation that will be processed under the assumption that it follows the role scheme for output relational-tuples specified in Sec. 4.1: the model needs to learn to produce TPRs such that this processing generates correct output programs. The decoding step is implemented by the TP-N2F relational tuples decoder (*TP-N2F Decoder*), which takes the output from the Reasoning Module (Sec. 4.1) and decodes the target sequence of relational tuples via TPR unbinding. The TP-N2F Decoder utilizes an attention mechanism over the individual-word TPRs \mathbf{T}^t produced by the TP-N2F Encoder. The detailed implementations are introduced below.

The TP-N2F encoder follows the role scheme in Sec. 4.1 to encode each word token w^t by soft-selecting one of n_F fillers and one of n_R roles. The fillers and roles are embedded as vectors. These embedding vectors, and the functions

for selecting fillers and roles, are learned by two LSTMs, the Filler-LSTM and the Role-LSTM. (See Figure 2.) At each time-step t , the Filler-LSTM and the Role-LSTM take a learned word-token embedding w^t as input. The hidden state of the Filler-LSTM, h_f^t , is used to compute softmax scores u_k^F over n_F filler slots, and a filler vector $f^t = F u^F$ is computed from the softmax scores (recall from Sec. 3 that F is the learned matrix of filler vectors). Similarly, a role vector is computed from the hidden state of the Role-LSTM, h_r^t . f_F and f_R denote the functions that generate f^t and r^t from the hidden states of the two LSTMs. The token w^t is encoded as \mathbf{T}^t , the tensor product of f^t and r^t . \mathbf{T}^t replaces the hidden vector in each LSTM and is passed to the next time step, together with the LSTM cell-state vector c^t : see (7)–(9). After encoding the whole sequence, the TP-N2F encoder outputs the sum of all tensor products $\sum_t \mathbf{T}^t$ to the next module. We use an MLP, called the Reasoning MLP, for TPR mapping; it takes an order-2 TPR from the encoder and maps it to the initial state of the decoder. Detailed equations and implementation are provided in Appendix.

$$h_F^t = f_{\text{Filler-LSTM}}(w^t, \mathbf{T}^{t-1}, c_F^{t-1}) \quad (7)$$

$$h_R^t = f_{\text{Role-LSTM}}(w^t, \mathbf{T}^{t-1}, c_R^{t-1}) \quad (8)$$

$$\mathbf{T}^t = f^t \otimes r^t = f_F(h_F^t) \otimes f_R(h_R^t) \quad (9)$$

Relational-Tuple Decoder in TP-N2F The TP-N2F Decoder is an RNN that takes the output from the reasoning MLP as its initial hidden state for generating a sequence of relational tuples (Figure 3). This decoder contains an attentional LSTM called the Tuple-LSTM which feeds an unbinding module: attention operates on the context vector of the encoder, consisting of all individual encoder outputs $\{\mathbf{T}^t\}$. The hidden-state \mathbf{H} of the Tuple-LSTM is treated as a TPR of a relational tuple and is unbound to a relation and

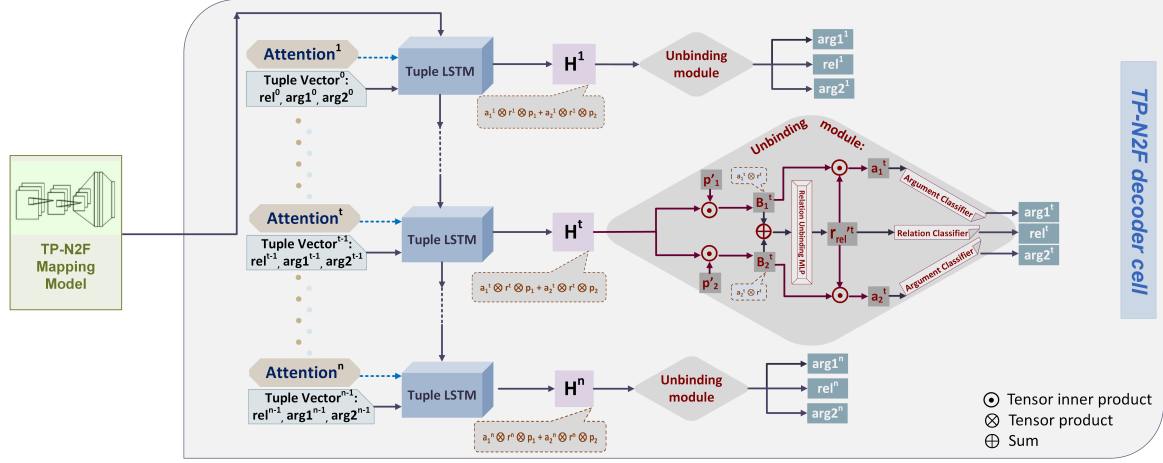


Figure 3. Implementation of the TP-N2F decoder.

arguments. During training, the Tuple-LSTM needs to learn a way to make \mathbf{H} suitably approximate a TPR. At each time step t , the hidden state \mathbf{H}^t of the Tuple-LSTM with attention (the version in Luong et al. (2015)) (11) is fed as input to the unbinding module, which regards \mathbf{H}^t as if it were the TPR of a relational tuple with m arguments possessing the role structure described in Sec. 4.1: $\mathbf{H}^t \approx \sum_{i=1}^m \mathbf{a}_i^t \otimes \mathbf{r}_{rel}^t \otimes \mathbf{p}_i$. (In Figure 3, the assumed hypothetical form of \mathbf{H}^t , as well as that of \mathbf{B}_i^t below, is shown in a bubble with dashed border.) To decode a binary relational tuple ($m = 2$), the unbinding module decodes it from \mathbf{H}^t using the two steps of TPR unbinding given in (4)–(5). The positional unbinding vectors \mathbf{p}_i^t are learned during training and shared across all time steps. After the first unbinding step (4), i.e., the inner product of \mathbf{H}^t with \mathbf{p}_i^t , we get tensors \mathbf{B}_i^t (12 and 13). These are treated as the TPRs of two arguments \mathbf{a}_i^t bound to a relation \mathbf{r}_{rel}^t . A relational unbinding vector \mathbf{r}_{rel}^{tt} is computed by a linear function from the sum of the \mathbf{B}_i^t (14) and used to compute the inner product with each \mathbf{B}_i^t to yield \mathbf{a}_i^t , which are treated as the embedding of argument vectors (15 and 16). Based on the TPR theory, \mathbf{r}_{rel}^{tt} is passed to a linear function to get \mathbf{r}_{rel}^t as the embedding of a relation vector. Finally, the softmax probability distribution over symbolic outputs is computed for relations and arguments separately. In generation, the most probable symbol is selected. (Detailed equations are in Appendix A.2.3 of supplementary.)

$$\mathbf{h}^t = f_{\text{Tuple-LSTM}}(\text{rel}^t, \text{arg}_1^t, \text{arg}_2^t, \mathbf{H}^{t-1}, \mathbf{c}^{t-1}) \quad (10)$$

$$\mathbf{H}^t = \text{Atten}(\mathbf{h}^t, [\mathbf{T}^0, \dots, \mathbf{T}^{n-1}]) \quad (11)$$

$$\mathbf{B}_1^t = \mathbf{H}^t \cdot \mathbf{p}'_1 \quad (12)$$

$$\mathbf{B}_2^t = \mathbf{H}^t \cdot \mathbf{p}'_2 \quad (13)$$

$$\mathbf{r}_{rel}^{tt} = f_{\text{linear}}(\mathbf{B}_1^t + \mathbf{B}_2^t) \quad (14)$$

$$\mathbf{a}_1^t = \mathbf{B}_1^t \cdot \mathbf{r}_{rel}^{tt} \quad (15)$$

$$\mathbf{a}_2^t = \mathbf{B}_2^t \cdot \mathbf{r}_{rel}^{tt} \quad (16)$$

4.3. Inference and Learning Strategy

During inference time, natural language questions are encoded via the encoder and the Reasoning MLP maps the output of the encoder to the input of the decoder. We use greedy decoding (selecting the most likely class) to decode one relation and its arguments. The relation and argument vectors are concatenated to construct a new vector as the input for the Tuple-LSTM in the next step.

TP-N2F is trained using back-propagation (Rumelhart et al., 1986) with the Adam optimizer (Kingma & Ba, 2017) and teacher-forcing. At each time step, the ground-truth relational tuple is provided as the input for the next time step. As the TP-N2F decoder decodes a relational tuple at each time step, the relation token is selected only from the relation vocabulary and the argument tokens from the argument vocabulary. For an input \mathcal{I} that generates N output relational tuples, the loss is the sum of the cross entropy loss \mathcal{L} between the true labels L and predicted tokens for relations and arguments as shown in (17).

$$\mathcal{L}_{\mathcal{I}} = \sum_{i=0}^{N-1} \mathcal{L}(\text{rel}^i, L_{\text{rel}^i}) + \sum_{i=0}^{N-1} \sum_{j=1}^2 \mathcal{L}(\text{arg}_j^i, L_{\text{arg}_j^i}) \quad (17)$$

4.4. Input of Decoder’s Unbinding Module is a TPR

Here we show that, if learning is successful, the order-3 tensor \mathbf{H} that each iteration of the decoder’s Tuple LSTM feeds to the decoder’s Unbinding Module (Figure 3) will be a TPR of the form assumed in Eq. 3 above, repeated here:

$$\mathbf{H} = \sum_{i \in I} \mathbf{a}_i \otimes \mathbf{r}_{rel} \otimes \mathbf{p}_i. \quad (18)$$

with $I \equiv \{1, \dots, \text{arity}(\text{rel})\}$. The decoder’s operations are given in Eqs.4–5, and Eqs.12–13, rewritten here:

$$\mathbf{H} \cdot \mathbf{p}'_i = \mathbf{q}_i \quad (19)$$

$$\mathbf{q}_i \cdot \mathbf{r}'_{rel} = \mathbf{a}_i \quad (20)$$

for $i \in I$. This is the standard TPR unbinding operation, used recursively: first with the unbinding vectors for positions, \mathbf{p}'_i , then with the unbinding vector for the operator, \mathbf{r}'_{rel} . It therefore suffices to analyze a single unbinding; the result can then be used recursively. This in effect reduces the problem to the order-2 case.

What we will show is: given a set of unbinding vectors $\{\mathbf{r}'_j\}$ which are dual to a set of role vectors $\{\mathbf{r}_j\}$, with j ranging over some index set J , if \mathbf{H} is an order-2 tensor such that

$$\mathbf{H} \cdot \mathbf{r}'_j = \mathbf{f}_j, \forall j \in J \quad (21)$$

then

$$\mathbf{H} = \sum_{j \in J} \mathbf{f}_j \mathbf{r}_j^\top + \mathbf{Z} \equiv \mathbf{H}_{\text{TPR}} + \mathbf{Z} \quad (22)$$

for some tensor \mathbf{Z} that annihilates all the unbinding vectors:

$$\mathbf{Z} \cdot \mathbf{r}'_j = \mathbf{0}, \forall j \in J. \quad (23)$$

Define “successful” learning to mean that the decoder’s processing will generate the target relational tuple (R, A_1, A_2) by obeying Eq. 19 in the first unbinding — which matches Eq. 21 with $\mathbf{r}'_j = \mathbf{p}'_j$, $\mathbf{f}_j = \mathbf{q}_j$, $J = I$ — and obeying Eq. 20 in the second unbinding — which, for each i , matches Eq. 21 with $\mathbf{H} = \mathbf{q}_i$, $\mathbf{r}'_j = \mathbf{r}'_{rel}$, $\mathbf{f}_j = \mathbf{a}_i$, $J = \{i\}$.

Treat rank-2 tensors as matrices; then unbinding is simply matrix-vector multiplication. Assume the set of unbinding vectors is linearly independent (otherwise there would not be a general way to satisfy Eq. 21 exactly, contrary to assumption). Then expand the set of unbinding vectors, if necessary, into a basis $\{\mathbf{r}'_k\}_{k \in K \supseteq J}$. Find the dual basis, with \mathbf{r}_k dual to \mathbf{r}'_k (so that $\mathbf{r}_l^\top \mathbf{r}'_j = \delta_{lj}$). Because $\{\mathbf{r}'_k\}_{k \in K}$ is a basis, so is $\{\mathbf{r}_k\}_{k \in K}$, so any matrix \mathbf{H} can be expanded as $\mathbf{H} = \sum_{k \in K} \mathbf{v}_k \mathbf{r}_k^\top$. Since $\mathbf{H} \mathbf{r}'_j = \mathbf{f}_j, \forall j \in J$ are the unbinding conditions (Eq. 21), we must have $\mathbf{v}_j = \mathbf{f}_j, \forall j \in J$. Let $\mathbf{H}_{\text{TPR}} \equiv \sum_{j \in J} \mathbf{f}_j \mathbf{r}_j^\top$. This is the desired TPR, with fillers \mathbf{f}_j bound to the role vectors \mathbf{r}_j which are the duals of the unbinding vectors \mathbf{r}'_j ($\forall j \in J$). Then we have $\mathbf{H} = \mathbf{H}_{\text{TPR}} + \mathbf{Z}$ (Eq. 22) where $\mathbf{Z} \equiv \sum_{k \in K, k \notin J} \mathbf{v}_k \mathbf{r}_k^\top$; so $\mathbf{Z} \mathbf{r}'_j = \mathbf{0}, \forall j \in J$ (Eq. 23). Thus, if training is successful, the model must have learned how to feed the decoder with order-3 TPRs with the structure posited in Eq. 18.

The argument so far addresses the case where the unbinding vectors are linearly independent, making it possible to satisfy Eq. 21 exactly. In relatively high-dimensional vector spaces, it will often happen that even when the number of

unbinding vectors exceeds the dimension of their space by a factor of 2 or 3 (which applies to the TP-N2F models presented here), there is a set of role vectors $\{\mathbf{r}_k\}_{k \in K}$ approximately dual to $\{\mathbf{r}'_k\}_{k \in K}$, such that, to a good approximation, $\mathbf{r}_l^\top \mathbf{r}'_k = \delta_{lk} \forall l, k \in K$. (Let the distribution of normalized unbinding vectors be approximately uniform on the unit sphere; then choosing the approximate dual vectors to equal the unbinding vectors will do, since they will be nearly orthonormal. If the $\{\mathbf{r}'_k\}_{k \in K}$ are not normalized, we rescale them, choosing $\mathbf{r}_k = \mathbf{r}'_k / \|\mathbf{r}'_k\|$.) When the number of such role vectors exceeds the dimension of the embedding space, they will be overcomplete, so while it is still true that any matrix \mathbf{H} can be expanded as above ($\mathbf{H} = \sum_{k \in K} \mathbf{v}_k \mathbf{r}_k^\top$), this expansion will no longer be unique. So while it remains true that \mathbf{H} a TPR, it is no longer uniquely decomposable into filler/role pairs. The claim above does not claim uniqueness in this sense, and remains true.)

5. Experiments

The proposed TP-N2F model is evaluated on two N2F tasks, generating operation sequences to solve math problems and generating Lisp programs. In both tasks, TP-N2F achieves state-of-the-art performance. We further analyze the behavior of the unbinding relation vectors in the proposed model. Results of each task and the analysis of the unbinding relation vectors are introduced in turn. Details of experiments and datasets are described in Appendix A.1 of the supplementary materials.

5.1. Generating Operations to Solve Math Problems

Given a natural-language math problem, we need to generate a sequence of operations (operators and corresponding arguments) from a set of operators and arguments to solve the given problem. Each operation is regarded as a relational tuple by viewing the operator as a relation, e.g., $(\text{add}, n1, n2)$. We test TP-N2F for this task on the MathQA dataset (Amini et al., 2019). The MathQA dataset consists of about 37k math word problems, each with a corresponding list of multi-choice options and the corresponding operation sequence. In this task, TP-N2F is deployed to generate the operation sequence given the question. The generated operations are executed with the execution script from Amini et al. (2019) to select a multi-choice answer. As there are about 30% noisy data (where the execution script returns the wrong answer when given the ground-truth program; see Appendix A.1 in the supplementary materials), we report both execution accuracy (of the final multi-choice answer after running the execution engine) and operation sequence accuracy (where the generated operation sequence must match the ground truth sequence exactly).

TP-N2F is compared to a baseline provided by the seq2prog model in Amini et al. (2019), an LSTM-based seq2seq

Table 1. Results on MathQA dataset testing set

MODEL	Operation Accuracy (%)	Execution Accuracy (%)
SEQ2PROG-orig	59.4	51.9
SEQ2PROG-best	66.97	54.0
LSTM2TP (ours)	68.21	54.61
TP2LSTM (ours)	68.84	54.61
TP-N2F (ours)	71.89	55.95

model with attention. Our model outperforms both the original seq2prog, designated SEQ2PROG-orig, and the best reimplemented seq2prog after an extensive hyperparameter search, designated SEQ2PROG-best. Table 1 presents the results. To verify the importance of the TP-N2F encoder and decoder, we conducted experiments to replace either the encoder with a standard LSTM (denoted LSTM2TP) or the decoder with a standard attentional LSTM (denoted TP2LSTM). We observe that both the TPR components of TP-N2F are important for achieving the observed performance gain relative to the baseline.

5.2. Generating Lisp Programs from NL Descriptions

Generating Lisp programs requires sensitivity to structural information because Lisp code and data can be regarded as tree-structured. Given a natural-language query, we need to generate code containing function calls with parameters. Each function call is a relational tuple, which has a function as the relation and parameters as arguments. We evaluate our model on the AlgoLisp dataset for this task and achieve state-of-the-art performance. The AlgoLisp dataset (Polosukhin & Skidanov, 2018) is a program synthesis dataset. Each sample contains a natural-language problem description, a corresponding Lisp program tree, and 10 input-output testing pairs. We parse the program tree into a straight-line sequence of tuples (same style as in MathQA). AlgoLisp provides an execution script to run the generated program and has three evaluation metrics: the accuracy of passing all test cases (Acc), the accuracy of passing 50% of test cases (50p-Acc), and the accuracy of generating an exactly matching program (M-Acc). AlgoLisp has about 10% noisy data (details in the Appendix), so we report results both on the full test set and the cleaned test set (in which all noisy testing samples are removed).

TP-N2F is compared with an LSTM seq2seq with attention model, the Seq2Tree model in Polosukhin & Skidanov (2018), and a seq2seq model with a pre-trained tree decoder from the Tree2Tree autoencoder (SAPS) reported in Bednarek et al. (2019). As shown in Table 2, TP-N2F outperforms all existing models on both the full test set and the cleaned test set. Ablation experiments with TP2LSTM and LSTM2TP show that, for this task, the TP-N2F Decoder is more helpful than TP-N2F Encoder. This may be

because Lisp code relies more heavily on structured representations. Comparing the generated programs, TP-N2F can generate longer programs than the LSTM-based Seq2Seq, e.g. TP-N2F correctly generates a program with 55 tuples but the LSTM-based Seq2Seq fails. Generated examples are presented in the Appendix.

6. Interpretation of Learned Structure

To interpret the structure learned by the model, we explore both the TP-N2F Encoder and the Decoder. For TP-N2F Encoder, the Softmax scores for roles and fillers of natural-language are analyzed on selected questions. We explore the significant fillers and roles of natural-language questions with large Softmax scores. Analysis shows that fillers tend to represent the semantic information and words or phrases with the same meaning tend to be assigned the same filler. For example, in the AlgoLisp dataset, “consider”, “you are given” and “given” are assigned to filler 146. “decrement”, “difference of” and “decremented by” are assigned to filler 43. “increment” and “add” are assigned to filler 105. In the MathQA dataset, “positive integer”, “positive number” and “positive digits” are assigned to filler 27. We also find that roles tend to represent the structured schemes of sentences. For example, Figure 4 shows the visualization of assigned roles for two different questions from the Algolisp dataset. Words with role 12 indicate the target of the questions to compute. Words with role 3 indicate required information to solve the questions. One interesting finding is that the second example from Figure 4 has two occurrences of “a” with different meanings. Therefore, although they are assigned the same role, they have different fillers. The detailed visualization of fillers is shown in Appendix.

For the TP-N2F Decoder, we extract the trained unbinding relation vectors and reduce the dimension of vectors via Principal Components Analysis. K-means clustering results on the average vectors are presented in Appendix A.6 of the supplementary material. Results show that unbinding vectors for operators or functions with similar semantics tend to be close to each other. For example, with 5 clusters in the MathQA dataset, arithmetic operators such as *add*, *subtract*, *multiply*, *divide* are clustered together, and operators related to *square* or *volume* of geometry are clustered together. With 4 clusters in the AlgoLisp dataset, par-

Table 2. Results of AlgoLisp dataset

MODEL	Full Testing Set			Cleaned Testing Set		
	Acc (%)	50p-Acc (%)	M-Acc (%)	Acc (%)	50p-Acc (%)	M-Acc (%)
Seq2Tree	61.0	—	—	—	—	—
LSTM2LSTM+atten	67.54	70.89	75.12	76.83	78.86	75.42
TP2LSTM (ours)	72.28	77.62	79.92	77.67	80.51	76.75
LSTM2TPR (ours)	75.31	79.26	83.05	84.44	86.13	83.43
SAPSpred-VH-Att-256	83.80	87.45	—	92.98	94.15	—
TP-N2F (ours)	84.02	88.01	93.06	93.48	94.64	92.78




Figure 4. Visualizations of selected roles in TP-N2F encoder for two questions in the AlgoLisp Dataset.

tial/lambda functions and sort functions are in one cluster, and string processing functions are clustered together. Note that there is no direct supervision to inform the model about the nature of the operations, and the TP-N2F decoder has induced this role structure using weak supervision signals from question/operation-sequence-answer pairs.

7. Conclusion and Future Work

In this paper we propose a new scheme for neural-symbolic relational representations and a new architecture, TP-N2F, for formal-language generation from natural-language descriptions. To our knowledge, TP-N2F is the first model that combines TPR binding and TPR unbinding in the encoder-decoder fashion. TP-N2F achieves the state-of-the-art on two instances of N2F tasks, showing significant structure

learning ability. The results show that both the TP-N2F encoder and the TP-N2F decoder are important for improving natural- to formal-language generation. By exploring the learned structures in both encoder and decoder, we show that TPRs enhance the interpretability of sequence-processing deep learning models and provide a step towards better understanding neural models. Next, we will combine large-scale deep learning models such as BERT with TP-N2F to take advantage of structure learning for other generation tasks.

Acknowledgements

We are grateful to Aida Amini from the University of Washington for providing execution scripts for the MathQA dataset.

References

- Amini, A., Gabriel, S., Lin, P., Kedzior, R. K., Choi, Y., and Hajishirzi, H. Mathqa: Towards interpretable math word problem solving with operation-based formalisms. In *NACCL*, 2019.
- Bednarek, J., Piaskowski, K., and Krawiec, K. Ain't nobody got time for coding: Structure-aware program synthesis from natural language. In *arXiv.org*, 2019.
- Cai, D. and Lam, W. Core semantic first: A top-down approach for amr parsing. In *arXiv:1909.04303*, 2019.
- Chen, K. and Forbus, K. D. Action recognition from skeleton data via analogical generalization over qualitative representations. In *Thirty-Second AAAI Conference*, 2018.
- Chen, K., Rabkina, I., McLure, M. D., and Forbus, K. D. Human-like sketch object recognition via analogical learning. In *Thirty-Third AAAI Conference*, volume 33, pp. 1336–1343, 2019.
- Crouse, M., McFate, C., and Forbus, K. D. Learning from unannotated qa pairs to analogically disambiguate and answer questions. In *Thirty-Second AAAI Conference*, 2018.
- Forbus, K., Liang, C., and Rabkina, I. Representation and computation in cognitive models. In *Top Cognitive System*, 2017.
- Gao, J., Galley, M., and Li, L. Neural approaches to conversational ai. *Foundations and Trends in Information Retrieval*, 13(2-3):127–298, 2019.
- Goldin-Meadow, S. and Gentner, D. *Language in mind: Advances in the study of language and thought*. MIT Press, 2003.
- Huang, Q., Smolensky, P., He, X., Wu, O., and Deng, L. Tensor product generation networks for deep nlp modeling. In *NAACL*, 2018.
- Huang, Q., Deng, L., Wu, D., Liu, c., and He, X. Attentive tensor product learning. In *Thirty-Third AAAI Conference*, volume 33, 2019.
- Kamath, A. and Das, R. A survey on semantic parsing. In *AKBC*, 2019.
- Kingma, D. P. and Ba, J. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2017.
- Langley, P. Crafting papers on machine learning. In Langley, P. (ed.), *Proceedings of the 17th International Conference on Machine Learning (ICML 2000)*, pp. 1207–1216, Stanford, CA, 2000. Morgan Kaufmann.
- Lee, K., Palangi, H., Chen, X., Hu, H., and Gao, J. Learning visual relation priors for image-text matching and image captioning with neural scene graph generators. *abs/1909.09953*, 2019. URL <http://arxiv.org/abs/1909.09953>.
- Lee, M., He, X., Yih, W.-t., Gao, J., Deng, L., and Smolensky, P. Reasoning in vector space: An exploratory study of question answering. In *ICLR*, 2016.
- Liao, Y., Bing, L., Li, P., Shi, S., Lam, W., and Zhang, T. Core semantic first: A top-down approach for amr parsing. In *EMNLP*, pp. 3855—3864, 2018.
- Luong, M.-T., Pham, H., and Manning, C. D. Effective approaches to attention-based neural machine translation. *EMNLP*, pp. 533–536, 2015.
- Martin, A. E. A compositional neural architecture for language. *Journal of Cognitive Neuroscience*, 2020.
- Palangi, H., Smolensky, P., He, X., and Deng, L. Question-answering with grammatically-interpretable representations. In *AAAI*, 2018.
- Polosukhin, I. and Skidanov, A. Neural program search: Solving programming tasks from description and examples. In *ICLR workshop*, 2018.
- Roads, B. and Love, B. Learning as the unsupervised alignment of conceptual systems. *arXiv preprint arXiv:1906.09012*, 2019.
- Rumelhart, D. E., Hinton, G. E., and Williams, R. J. Learning internal representations by error propagation. In Rumelhart, D. E., McClelland, J. L., and the PDP Group (eds.), *Parallel distributed processing: Explorations in the microstructure of cognition*, volume 1, pp. 318–362. MIT press, Cambridge, MA, 1986.
- Schlag, I. and Schmidhuber, J. Learning to reason with third order tensor products. In *Neural Information Processing Systems*, 2018.
- Smolensky, P. Tensor product variable binding and the representation of symbolic structures in connectionist networks. In *Artificial Intelligence*, volume 46, pp. 159–216, 1990.
- Smolensky, P., Lee, M., He, X., Yih, W.-t., Gao, J., and Deng, L. Basic reasoning with tensor product representations. *arXiv preprint arXiv:1601.02745*, 2016.