

Neural Programming by Example

Chengxun Shu

Beihang University
Beijing 100191, China
shuchengxun@163.com

Hongyu Zhang

The University of Newcastle
Callaghan, NSW 2308, Australia
hongyu.zhang@newcastle.edu.au

Abstract

Programming by Example (PBE) targets at automatically inferring a computer program for accomplishing a certain task from sample input and output. In this paper, we propose a deep neural networks (DNN) based PBE model called Neural Programming by Example (NPBE), which can learn from input-output strings and induce programs that solve the string manipulation problems. Our NPBE model has four neural network based components: a string encoder, an input-output analyzer, a program generator, and a symbol selector. We demonstrate the effectiveness of NPBE by training it end-to-end to solve some common string manipulation problems in spreadsheet systems. The results show that our model can induce string manipulation programs effectively. Our work is one step towards teaching DNN to generate computer programs.

Introduction

Programming by Example (PBE, also called programming by demonstration, or inductive synthesis) (Lieberman 2001; Cypher and Halbert 1993; Gulwani 2011) gives machines the ability to reason and generate new programs without substantial amount of human supervision. In PBE systems, users (often non-professional programmers) provide a machine with input-output examples of a task they would like to perform and the machine automatically infers a program to accomplish the task. The concept of PBE has been successfully used for string manipulation in spreadsheet systems such as Microsoft Excel (Gulwani et al. 2015), computer-aided education (Gulwani 2014) and data extracting systems (Le and Gulwani 2014).

As an example, if a user provides the following input and output examples:

$john@example.com \Rightarrow john$
 $james@company.com \Rightarrow james$

A PBE system should understand that the user would like to extract the user name from the email address. It will automatically synthesize a program $Select(Split(x, @), 0)$, where x is the input string, $Split$ is to split a string according to a delimiter, and $Select$ is to select a substring from an array

of strings. Given a new email address $jacob@test.com$, the program will output the string $jacob$.

Lau et al. (2003) applied version space algebra to search for possible programs. More recent PBE methods (Gulwani 2011) mainly adopt search technique to find a composition of predefined functions (such as string Split and Concatenation) that satisfies the input-output examples. These methods create a Directed Acyclic Graph (DAG) and search through the sequences of functions that can generate the output string from a given input state. These methods can generate complex string manipulation programs effectively, but require the design of complex program synthesis algorithms.

In this paper, we propose a Deep Neural Networks (DNN) based approach to Programming by Example. We train neural networks to automatically infer programs from input-output examples. During the past few years, research on DNN has achieved significant results in a variety of fields such as computer vision (Krizhevsky, Sutskever, and Hinton 2012), speech recognition (Mohamed, Dahl, and Hinton 2012), natural language processing (Collobert et al. 2011), and API learning (Gu et al. 2016). Recently, researchers have explored the feasibility of applying DNN to solve programming and computation related problems (Neelakantan, Le, and Sutskever 2015; Reed and de Freitas 2015; Graves, Wayne, and Danihelka 2014; Kurach, Andrychowicz, and Sutskever 2015). Our work is based on the similar idea of applying DNN to infer and execute computer programs. Different from the existing work, we target at the PBE problem and train a neural PBE model with triples of input, output and program.

Our approach, called NPBE (Neural Programming by Example), teaches DNN to compose a set of predefined atomic operations for string manipulation. Given an input-output string pair, the NPBE model is trained to synthesize a program - a sequence of functions and corresponding arguments that transform the input string to the output string. The program is generated from the atomic functions and one function may use the execution results of previous functions. Thus the model is able to compose complex programs using only several predefined operations.

We have experimentally evaluated NPBE on a large number of input-output strings for 45 string manipulation tasks and the results are encouraging. We find that our model can generalize beyond the training input/output strings and the

training argument settings. Our work is one of the early attempts to apply DNN to the problem of Programming by Example.

Related Work

Reed et al. (2015) developed a framework called Neural Programmer-Interpreter to induce and execute programs using neural networks. Their method treats programs as embeddings and uses neural networks to generate functions and arguments for program execution. Their model is trained with the supervision of execution traces and can be used in many different scenarios. However, their model cannot be directly applied to the problem of PBE as the input to their model is the environment encoded with the task, while our work is dedicated to PBE and the input to our model are input-output examples. Neural Programmer (Neelakantan, Le, and Sutskever 2015) is a neural network augmented with a set of operations that can be called over several steps. It is trained to output the result of program execution, while our model is trained to output the program represented by symbols. Neural Enquirer (Yin et al. 2015) is a fully neural, end-to-end differentiable model capable of modeling and executing table query related programs. The execution of Neural Enquirer is “softly” on data table using neural networks, while our work does not apply soft execution to input-output strings. Bošnjak et al. (2016) proposed a neural implementation of an abstract machine for the language Forth. It can learn program behaviour trained from input-output data. However, it requires program sketches as input.

Our model is also related to the work that uses recurrent neural networks to solve programming and computation related problems. Graves et al. (2014) developed Neural Turing Machine which is capable of learning and executing simple programs using an external memory. Zaremba et al. (2015) used execution traces to train recurrent neural networks to learn simple algorithms. Ling et al. (2016) developed a model to generate program code from natural language and structured specification. Pointer Networks (Vinyals, Fortunato, and Jaitly 2015) uses an attentional recurrent model to solve difficult algorithmic problems.

Some machine learning methods were also proposed to tackle PBE problems. Lau et al. (2003) applied version space algebra to efficiently search for possible programs. Given input-output pairs, genetic programming (Banzhaf et al. 1998) can evolve useful programs from candidate populations. Liang et al. (2010) proposed a hierarchical Bayesian approach to learn simple programs given only a few examples. Melon et al. (2013) used machine learning to speed up the searching for possible programs by learning weights related to textual features. Their method needs carefully designed features to reduce search space, while our method reduces search space and avoids feature engineering through learning representations using DNN.

The NPBE Model

Problem Statement: Let S denote the set of strings. For an input string $x \in S$ and an output string $y \in S$, our model is to induce a function $f \in S^S$, so that $y = f(x)$. The input

to our model is the input-output pair $z := (x, y) \in S^2$, and we wish to get the correct function f , so when the user input another string $\bar{x} \in S$, we can generate the desired output $\bar{y} = f(\bar{x})$ without the need of user specifying the function explicitly.

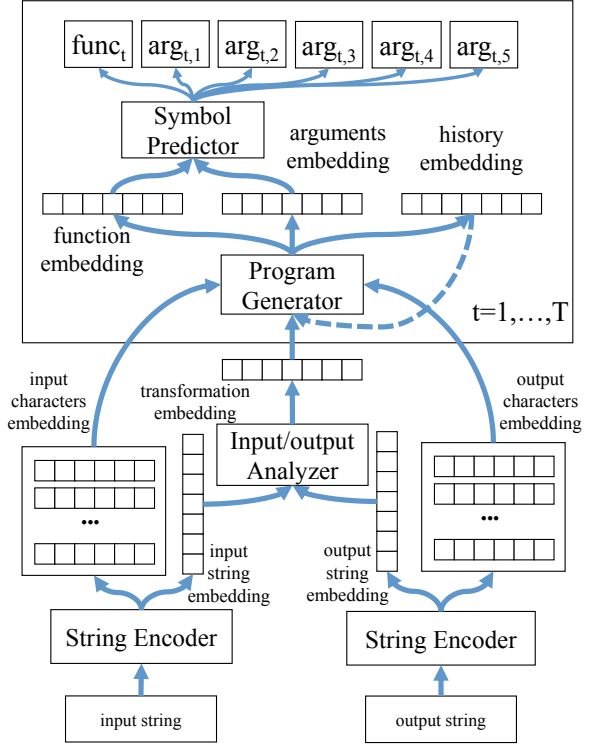


Figure 1: Overall architecture of NPBE.

The proposed NPBE model (Figure 1) consists of four modules:

- A string encoder to encode input and output strings;
- An input-output analyzer which generates the transformation embedding describing the relationship between input and output strings;
- A program generator which produces the function/arguments embeddings over a few steps;
- A symbol selector to decode the function/arguments embeddings and generate human readable symbols.

The program generator will run for a few steps, each step it may access the outputs of previous steps, thus enables the model to compose powerful programs using only a few predefined atomic functions. While traditional deep learning models try to generate the output y given x and eventually fit the function f such that $y = f(x)$, our model learns to generate the function f directly and fits the higher-order function g for which $f = g(x, y)$ and f satisfies $y = f(x)$.

String Encoder

Given a string composed of a sequence of characters $\{i_1, i_2, \dots, i_L\}$, the string encoder outputs a string embed-

ding and a list of character embeddings. First each character i_k in the sequence is mapped to the 8-dimensional raw embeddings $e(i_k)$ via a randomly initialized and trainable embedding matrix. To better present and attend to a character, the context of each character is fused with the character's raw embedding to build the character embedding. Let $l(i_k)$ denote the left context of the character i_k and $r(i_k)$ as the right context of i_k . $l(i_k)$ and $r(i_k)$ are respectively calculated using Equation (1) and Equation (2) with $l(0) = r(L+1) = [0]^C$. In our implementation f_l and f_r are the update function of LSTM (Hochreiter and Schmidhuber 1997).

$$l(i_k) = f_l(l(i_{k-1}), e(i_k)) \quad (1)$$

$$r(i_k) = f_r(r(i_{k+1}), e(i_k)) \quad (2)$$

The character embedding for each character is the combination of left/right contexts of character i_k and $e(i_k)$ itself as shown in Equation (3), where $[\cdot; \cdot]$ means the concatenation of vectors, \max is the element-wise max-pooling. W_e and b_e are parameter matrix and vector for building the full character embedding.

$$c_k = \tanh(W_e[\max(l(i_k), r(i_k)); e(i_k)] + b_e) \quad (3)$$

c_k will be used by the attention mechanism in the program generator.

We also need a representation which can summarize the whole string, so we can induce the transformation from the string embeddings of input and output strings. We use multilayer bidirectional LSTM (Graves and Schmidhuber 2005) to summarize c_k . The output of forward and backward LSTM at each layer are concatenated and become the input of next layer's forward and backward LSTM. The top-most layer's last hidden states of forward and backward LSTM are merged to generate the string embedding $s \in R^S$ through a final fully connected layer. The processing for input and output strings is separated but shares the same neural network architecture and parameters, thus producing $c_{I,1}, \dots, c_{I,L}, s_I$ and $c_{O,1}, \dots, c_{O,L}, s_O$ for input and output strings, respectively.

Input/output Analyzer

The input/output analyzer converts the input and output string embeddings to the transformation embedding, which describes the relationship between the input and output strings. Let $t \in R^T$ denote the transformation embedding. $s_I \in R^S, s_O \in R^S$ are the input and output string embeddings, respectively. The input/output analyzer can be represented as Equation (4). In our implementation, f_{IO} is just a 2-layer fully connected neural network with activation function \tanh .

$$t = f_{IO}([s_I; s_O]) \quad (4)$$

Program Generator

The program generator (Figure 2) is the core of NPBE. It generates several functions and corresponding arguments' embeddings step by step. The function's embedding at time t is calculated as Equation (5), which is a fully connected

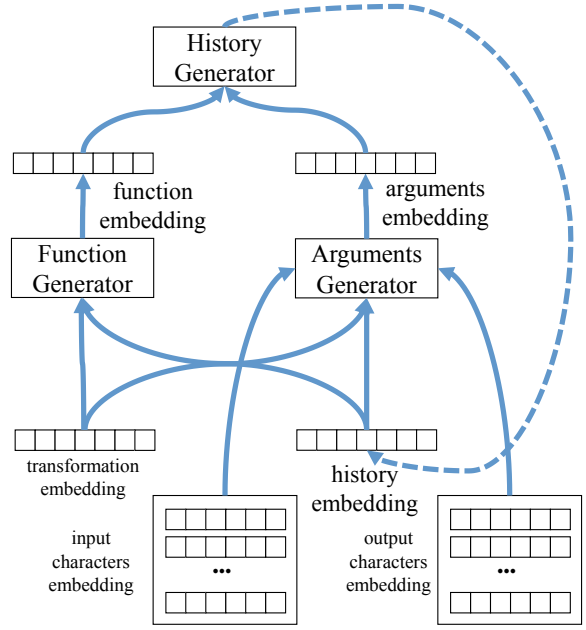


Figure 2: Diagram of the program generator.

neural network taking as input the transformation embedding t and the execution history of the program generator $h_{t-1} \in R^H$ ($h_0 = t$). Similarly, the function's arguments embedding at time t is calculated as Equation (6). However, the function's arguments are often very complex and hard to predict accurately. So attention mechanism (similar to the one used in neural machine translation models (Bahdanau, Cho, and Bengio 2014)) is applied to refine the raw arguments embedding $a_{r,t}$ with attention on input and output strings. This is summarized in Equation (7). Finally, the function embedding f_t , the refined arguments embedding a_t , and the previous history embedding h_{t-1} are merged into the new history embedding h_t as shown in Equation (8). W_h and b_h are parameter matrix and vector for generating the new history, respectively.

$$f_t = f_{func}(t, h_{t-1}) \quad (5)$$

$$a_{r,t} = f_{args}(t, h_{t-1}) \quad (6)$$

$$a_t = f_{att}(a_{r,t}, c_{I,1}, \dots, c_{I,L}, c_{O,1}, \dots, c_{O,L}) \quad (7)$$

$$h_t = \tanh(W_h[f_t; a_t; h_{t-1}] + b_h) \quad (8)$$

Functions f_{func} and f_{args} can be multilayer fully connected neural networks, and in our experiment we just use one layer neural network. The function f_{att} in Equation (7) is implemented by attending to the input and output strings as follows:

Input-output Pair	t	Function	Argument 1	Argument 2	Argument 3	Argument 4	Argument 5
Input string:	1	GetConstString	—	—	—	—	—
john@company.com	2	GetConstString	—	—	—	—	—
Output string:	3	Split	x	“@”	—	—	—
Hello john, have fun!	4	Select	o3	0	—	—	—
	5	Concatenate	o1	o4	o2	—	—
Input string:	1	Split	x	“/”	—	—	—
17/apr/2016	2	Select	o1	0	—	—	—
Output string:	3	Select	o1	1	—	—	—
APR-17	4	ToUpper	o3	—	—	—	—
	5	Concatenate	o4	“-”	o2	—	—
Input string:	1	Split	x	“/”	—	—	—
/home/foo/file.cpp	2	Select	o1	-1	—	—	—
Output string:	3	Split	o2	“.”	—	—	—
file	4	Select	o3	0	—	—	—

Table 3: Examples of input-output pairs for NPBE. The corresponding programs (from top to bottom) are: 1) *Concatenate*(“Hello ”, *Select*(*Split*(x, “@”), 0), “, have fun!”); 2) *Concatenate*(*ToUpper*(*Select*(*Split*(x, “/”), 1)), “-”, *Select*(*Split*(x, “/”), 0)); and 3) *Select*(*Split*(*Select*(*Split*(x, “/”), -1), “.”), 0). Note that *GetConstString* is a placeholder, and the first one is evaluated to “Hello”, the second one is evaluated to “, have fun!”.

string or join an array of strings by a delimiter. We also define some special symbols. For example, the symbol x refers to the input string, $o1$, $o2$, $o3$ and $o4$ are used to refer to the output of the first, second, third and fourth operation, respectively. The *NoArg* symbol indicates that no arguments is expected at the current position. Note that although in our experiments, we give some constraints to the functions and arguments in a program, our model can be easily extended to support new functions and arguments.

To obtain training data, we first generate programs at various levels of complexity according to 45 predefined tasks. A task is a sequence of function in a specific order, but the arguments of each function are not fixed. The reason for defining tasks is that we want the program generated by our model being syntactically correct and meaningful. The 45 tasks range from simple ones such as the concatenation of input to some constant string, to more complex ones comprising *Split*, *Join*, *Select*, *ToUpper* and *Concatenate* functions. For example, the task of *Split*, *Join* is to first split the input string by a delimiter, then join the split strings array using another delimiter. A program derived from this task could be *Join*(*Split*(x, “/”), “.”), which splits the input string x according to the delimiter “/” and then joins the resulting substrings using the delimiter “.”. The average number of functions for accomplishing a task is 3.5.

For all the tasks, we generate a total number of around 69,000 programs for training. For each program \mathcal{P}_i we generate a random input string \mathcal{X}_i , which should be an valid input to the program \mathcal{P}_i . Next, we apply the program \mathcal{P}_i on \mathcal{X}_i by actually running the Python program implementing \mathcal{P}_i and get the output string \mathcal{Y}_i . We constrain \mathcal{X}_i and \mathcal{Y}_i to be at most 62 characters long. After that, we use $\{\mathcal{X}_i, \mathcal{Y}_i\}$ as the input data to our model, and \mathcal{P}_i as the training target. The program is generated in such a way that if there are multiple programs that can result in the same \mathcal{Y}_i given \mathcal{X}_i , only one specific program is chosen. Therefore, there is no ambiguity for the model to predict the desired program. Given the program \mathcal{P}_i , the input string \mathcal{X}_i is always generated

dynamically and randomly to decrease overfitting. Table 3 gives some concrete input-output examples for our model.

To train NPBE, we choose RMSProp (Tieleman and Hinton 2012) as the optimizer and set the mini-batch size to 200. We set the dimensionality of the transformation embedding \mathbf{t} and the history embedding \mathbf{h} to 256, the function embedding \mathbf{f} to 16, and the arguments embedding \mathbf{a} to 64. The actual training process relies on an adaptive curriculum (Reed and de Freitas 2015) in which the frequency of one specific task being trained is proportional to its error rates over test. Every 10 epochs we estimate the prediction errors. We use softmax with adequate temperature over error rates to sample the frequency of each task that will be trained during the next 10 epochs. So tasks with the higher error rates will be sampled more frequently than tasks with the lower error rates during the next 10 epochs.

The evaluation of NPBE is conducted to answer the following research questions:

RQ1: What is the accuracy of NPBE in generating programs? In this RQ, we use randomly generated input-output strings to evaluate the accuracy of NPBE in generating programs. For example, given a random input-output pair: 25/11/16 and 25:11:16 (which does not appear in the training data), we would like to test if the correct program *Join*(*Split*(x, “/”), “.”) can be still generated. To answer this RQ, we generate random input-output strings for each task 1000 times and apply the trained NPBE model. A program produced by NPBE is regarded correct only if the model predicts all the five functions (if less than five, padded with the *NoFunc* symbol) and all the arguments of the functions (also padded with the *NoArg* symbol) correctly (thus a total of $5 + 5 \times 5 = 30$ positions). We also compare our model with the RNN encoder-decoder models (Cho et al. 2014) implemented using LSTM and LSTM with attention mechanism, which all have similar total number of parameters to our model.

RQ2: Can NPBE generate programs with previously unseen arguments? In this RQ, we test the generalization

Task	LSTM	LSTM-A	NPBE-Top1	NPBE-Top3	NPBE-Top5
Case Change	100.0%	100.0%	100.0%	100.0%	100.0%
Duplicate Input String	9.4%	9.8%	99.1%	99.1%	99.1%
Case Change and Concatenate with Input String	9.4%	9.3%	99.9%	99.9%	99.9%
Concatenate with Constant	83.1%	67.5%	88.4%	88.4%	88.4%
Split, Select, Case Change	8.7%	12.8%	92.3%	96.4%	97.3%
Split, Join, Select, Case Change, Concatenate	0.1%	0.1%	92.4%	96.6%	97.8%
GetConstString, Split, Select, Case Change, Concatenate	2.9%	5.2%	81.2%	85.6%	86.7%
GetConstString \times 2, Split, Select, Concatenate	0.1%	0.1%	24.7%	54.4%	68.1%
Split, Select, Select, Concatenate	0.1%	0.1%	9.8%	46.4%	82.0%
Split, Select, Select, Case Change, Concatenate	0.1%	0.2%	35.8%	73.2%	94.0%
Split, Select, Split, Select	0.1%	0.1%	31.2%	64.3%	72.2%
Average over All 45 Tasks	20.4%	22.0%	74.1%	85.8%	91.0%

Table 4: Results of generating programs.

Task	Seen	Unseen
Split, Join	93.6%	93.9%
GetConstString, Split, Join, Concatenate	91.4%	91.5%
Split, Join, Concatenate	95.0%	94.9%
Split, Join, Select, Concatenate	90.8%	89.5%
Split, Select, Select, Concatenate	82.0%	82.1%
Average over 19 Tasks	87.6%	87.4%

Table 5: Results with seen and unseen program arguments.

ability of NPBE. We evaluate our model using programs whose argument settings do not appear in the training set. For example, if the program *Join(Split(x, "/"), ":",")* appears in the training set, we would like to know if NPBE can work for the program *Join(Split(x, "@"), "-")*, which does not appear in the training set. To answer this RQ, we design a test set consisting of around 19,000 programs with previously unseen arguments. The experiment is conducted on 19 selected tasks that have complex argument combinations (for simple tasks there are few arguments to choose so we skip them).

Experimental Results

RQ1: What is the accuracy of NPBE in generating programs? Table 4 gives the evaluation results of NPBE on predicting programs. The average Top1 accuracy achieved by NPBE is 74.1%, which means that for 74.1% of input-output pairs in test, NPBE successfully generates the corresponding program. We found that the model prediction errors most likely to occur on the integer argument of *Select* because neural networks are not good at counting. So we also let the model to give 3 or 5 predictions when it tries to predict the integer arguments of *Select*. The average Top3 and Top5 accuracy results are 85.8% and 91.0%, which means that for 85.8% and 91.0% of input-output pairs in test, NPBE successfully returns the corresponding program within the top 3 and top 5 results respectively. The results show that the NPBE model can generate correct programs for most tasks.

As an example, given the input string "17/apr/2016" and output string "APR-17", our model needs to induce a program comprising *Split, Select, Select, Case Change, Concatenate*. For this task, our model gives completely correct

program in 35.8% cases. If we allow the model to give 3 or 5 predictions for the integer argument of *Select*, the accuracy is increased to 73.2% or 94.0%.

The results about LSTM and LSTM with attention mechanism (denoted to as LSTM-A) are also shown in Table 4. Note that the Top1, Top3 and Top5 accuracy for LSTM and LSTM-A are almost the same, so only the Top1 accuracy is reported. We found that the ordinary encoder-decoder model can solve the simplest tasks but cannot tackle harder tasks. The results show that NPBE significantly outperforms the ordinary encoder-decoder model.

RQ2: Can NPBE generate programs with previously unseen arguments? We test the generalization ability of NPBE on the programs with previously unseen argument settings. The average Top5 accuracy results are given in Table 5. The results shows that for seen and unseen argument settings the accuracies achieved by our model have no big difference. For example, the task of *Split, Join* first splits an input string by a delimiter (such as "/") and then concatenates the resulting substrings by the other delimiter (such as ":"). This task achieves 93.6% accuracy on the training set. For different arguments (e.g., first split the input string by "@" then join the resulting substrings by "-") that do not exist in the training set, NPBE can still get 93.9% accuracy. The results show that NPBE can be generalized to unseen program arguments without over-fitting to particular argument combinations.

Discussions and Future Work

The intention behind NPBE is to make the model learn related features from input-output strings automatically and use the learned features to induce correct programs. The purpose of this paper is not to directly compete with the existing PBE systems. Instead, we show that the use of DNN can recognize features in string transformations and can learn accurate programs through input-output pairs.

Currently, NPBE cannot be generalized to completely unseen tasks (such as *Split, Join, Join, Concatenate*) that never appeared in the training set. In our future work, we will try to build the model that really "understands" the meaning of atomic functions to make it possible to generalize to the unseen tasks.

Conclusion

In this paper, we propose NPBE, a Programming by Example (PBE) model based on DNN. NPBE can induce string manipulation programs based on simple input-output pairs by inferring a composition of functions and corresponding arguments. We have shown that the novel use of DNN can be successfully applied to develop Programming By Example systems. Our work also explores the way of learning higher-order functions in deep learning, and is one step towards teaching DNN to generate computer programs.

References

- Bahdanau, D.; Cho, K.; and Bengio, Y. 2014. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*.
- Banzhaf, W.; Nordin, P.; Keller, R. E.; and Francone, F. D. 1998. *Genetic programming: an introduction*. Morgan Kaufmann Publishers San Francisco.
- Bošnjak, M.; Rocktäschel, T.; Naradowsky, J.; and Riedel, S. 2016. Programming with a differentiable forth interpreter. *arXiv preprint arXiv:1605.06640*.
- Cho, K.; Merriënboer, B. V.; Gulcehre, C.; Bahdanau, D.; Bougares, F.; Schwenk, H.; and Bengio, Y. 2014. Learning phrase representations using rnn encoder-decoder for statistical machine translation. *arXiv preprint arXiv:1406.1078*.
- Collobert, R.; Weston, J.; Bottou, L.; Karlen, M.; Kavukcuoglu, K.; and Kuksa, P. 2011. Natural language processing (almost) from scratch. *Journal of Machine Learning Research* 12(Aug):2493–2537.
- Cypher, A., and Halbert, D. C. 1993. *Watch what I do: programming by demonstration*. MIT press.
- Graves, A., and Schmidhuber, J. 2005. Framewise phoneme classification with bidirectional lstm and other neural network architectures. *Neural Networks* 18(5-6):602–10.
- Graves, A.; Wayne, G.; and Danihelka, I. 2014. Neural turing machines. *arXiv preprint arXiv:1410.5401*.
- Gu, X.; Zhang, H.; Zhang, D.; and Kim, S. 2016. Deep API learning. In *Proceedings of the 24th ACM SIGSOFT International Symposium on Foundations of Software Engineering, FSE 2016*, 631–642. New York, NY, USA: ACM.
- Gulwani, S.; Hernandez-Orallo, J.; Kitzelmann, E.; Muggleton, S. H.; Schmid, U.; and Zorn, B. 2015. Inductive programming meets the real world. *Communications of the ACM* 58(11):90–99.
- Gulwani, S. 2011. Automating string processing in spreadsheets using input-output examples. In *ACM SIGPLAN Notices*, volume 46, 317–330. ACM.
- Gulwani, S. 2014. Example-based learning in computer-aided stem education. *Communications of the ACM* 57(8):70–80.
- Hochreiter, S., and Schmidhuber, J. 1997. Long short-term memory. *Neural Computation* 9(8):1735–1780.
- Krizhevsky, A.; Sutskever, I.; and Hinton, G. E. 2012. Imagenet classification with deep convolutional neural networks. In *Advances in neural information processing systems*, 1097–1105.
- Kurach, K.; Andrychowicz, M.; and Sutskever, I. 2015. Neural random-access machines. *arXiv preprint arXiv:1511.06392*.
- Lau, T.; Wolfman, S. A.; Domingos, P.; and Weld, D. S. 2003. Programming by demonstration using version space algebra. *Machine Learning* 53(1-2):111–156.
- Le, V., and Gulwani, S. 2014. Flashextract: a framework for data extraction by examples. In *ACM SIGPLAN Notices*, volume 49, 542–553. ACM.
- Liang, P.; Jordan, M. I.; and Klein, D. 2010. Learning programs: A hierarchical bayesian approach. In *Proceedings of the 27th International Conference on Machine Learning (ICML-10)*, 639–646.
- Lieberman, H. 2001. *Your wish is my command: Programming by example*. Morgan Kaufmann.
- Ling, W.; Grefenstette, E.; Hermann, K. M.; Kocisky, T.; Senior, A.; Wang, F.; and Blunsom, P. 2016. Latent predictor networks for code generation. *arXiv preprint arXiv:1603.06744*.
- Menon, A. K.; Tamuz, O.; Gulwani, S.; Lampson, B. W.; and Kalai, A. 2013. A machine learning framework for programming by example. *ICML (1)* 28:187–195.
- Mohamed, A.-r.; Dahl, G. E.; and Hinton, G. 2012. Acoustic modeling using deep belief networks. *IEEE Transactions on Audio, Speech, and Language Processing* 20(1):14–22.
- Neelakantan, A.; Vilnis, L.; Le, Q. V.; Sutskever, I.; Kaiser, L.; Kurach, K.; and Martens, J. 2015. Adding gradient noise improves learning for very deep networks. *arXiv preprint arXiv:1511.06807*.
- Neelakantan, A.; Le, Q. V.; and Sutskever, I. 2015. Neural programmer: Inducing latent programs with gradient descent. *arXiv preprint arXiv:1511.04834*.
- Reed, S., and de Freitas, N. 2015. Neural programmer-interpreters. *arXiv preprint arXiv:1511.06279*.
- Tieleman, T., and Hinton, G. 2012. Lecture 6.5-rmsprop: Divide the gradient by a running average of its recent magnitude. *COURSERA: Neural Networks for Machine Learning* 4(2).
- Vinyals, O.; Fortunato, M.; and Jaitly, N. 2015. Pointer networks. In *Advances in Neural Information Processing Systems*, 2692–2700.
- Yin, P.; Lu, Z.; Li, H.; and Kao, B. 2015. Neural enquirer: Learning to query tables with natural. *arXiv preprint arXiv:1512.00965*.
- Zaremba, W.; Mikolov, T.; Joulin, A.; and Fergus, R. 2015. Learning simple algorithms from examples. *arXiv preprint arXiv:1511.07275*.