

Easy over Hard: A Case Study on Deep Learning

Wei Fu, Tim Menzies

Com.Sci., NC State, USA

wfu@ncsu.edu, tim.menzies@gmail.com

ABSTRACT

While deep learning is an exciting new technique, the benefits of this method need to be assessed with respect to its computational cost. This is particularly important for deep learning since these learners need hours (to weeks) to train the model. Such long training time limits the ability of (a) a researcher to test the stability of their conclusion via repeated runs with different random seeds; and (b) other researchers to repeat, improve, or even refute that original work.

For example, recently, deep learning was used to find which questions in the Stack Overflow programmer discussion forum can be linked together. That deep learning system took 14 hours to execute. We show here that applying a very simple optimizer called DE to fine tune SVM, it can achieve similar (and sometimes better) results. The DE approach terminated in 10 minutes; i.e. 84 times faster hours than deep learning method.

We offer these results as a cautionary tale to the software analytics community and suggest that not every new innovation should be applied without critical analysis. If researchers deploy some new and expensive process, that work should be baselined against some simpler and faster alternatives.

CCS CONCEPTS

• **Computing methodologies** → **Neural networks**; • **Theory of computation** → **Evolutionary algorithms**; • **Software and its engineering** → **Search-based software engineering**; **Collaboration in software development**;

KEYWORDS

Search based software engineering, software analytics, parameter tuning, data analytics for software engineering, deep learning, SVM, differential evolution

ACM Reference Format:

Wei Fu, Tim Menzies. 2017. Easy over Hard: A Case Study on Deep Learning. In *Proceedings of 2017 11th Joint Meeting of the European Software Engineering Conference and the ACM SIGSOFT Symposium on the Foundations of Software Engineering, Paderborn, Germany, September 4–8, 2017 (ESEC/FSE’17)*, 12 pages.
<https://doi.org/10.1145/3106237.3106256>

Permission to make digital or hard copies of all or part of this work for personal or classroom use is granted without fee provided that copies are not made or distributed for profit or commercial advantage and that copies bear this notice and the full citation on the first page. Copyrights for components of this work owned by others than ACM must be honored. Abstracting with credit is permitted. To copy otherwise, or republish, to post on servers or to redistribute to lists, requires prior specific permission and/or a fee. Request permissions from permissions@acm.org.

ESEC/FSE’17, September 4–8, 2017, Paderborn, Germany

© 2017 Association for Computing Machinery.

ACM ISBN 978-1-4503-5105-8/17/09...\$15.00

<https://doi.org/10.1145/3106237.3106256>

1 INTRODUCTION

This paper extends a prior result from ASE’16 by Xu et al. [74] (hereafter, XU). XU described a method to explore large programmer discussion forums, then uncover related, but separate, entries. This is an important problem. Modern SE is evolving so fast that these forums contain more relevant and recent comments on current technologies than any textbook or research article.

In their work, XU predicted whether two questions posted on Stack Overflow are semantically linkable. Specifically, XU define a question along with its entire set of answers posted on Stack Overflow as a *knowledge unit* (KU). If two knowledge units are semantically related, they are considered as *linkable* knowledge units.

In their paper, they used a convolution neural network (CNN), a kind of deep learning method [42], to predict whether two KUs are linkable. Such CNNs are highly computationally expensive, often requiring network composed of 10 to 20 layers, hundreds of millions of weights and billions of connections between units [42]. Even with advanced hardware and algorithm parallelization, training deep learning models still requires hours to weeks. For example:

- XU report that their analysis required 14 hours of CPU.
- Le [40] used a cluster with 1,000 machines (16,000 cores) for three days to train a deep learner.

This paper debates what methods should be recommended to those wishing to repeat the analysis of XU. We focus on whether using simple and faster methods can achieve the results that are currently achievable by the state-of-art deep learning method. Specifically, we repeat XU’s study using DE (differential evolution [62]), which serves as a hyper-parameter optimizer to tune XU’s baseline method, which is a conventional machine learning algorithm, support vector machine (SVM). Our study asks:

RQ1: *Can we reproduce XU’s baseline results (Word Embedding + SVM)?* Using such a baseline, we can compare our methods to those of XU.

RQ2: *Can DE tune a standard learner such that it outperforms XU’s deep learning method?* We apply differential evolution to tune SVM. In terms of precision, recall and F1-score, we observe that the tuned SVM method outperforms CNN in most evaluation scores.

RQ3: *Is tuning SVM with DE faster than XU’s deep learning method?* Our DE method is 84 times faster than CNN.

We offer these results as a cautionary tale to the software analytics community. While deep learning is an exciting new technique, the benefits of this method need to be carefully assessed with respect to its computational cost. More generally, if researchers deploy some new and expensive process (like deep learning), that work should be baselined against some simpler and faster alternatives

The rest of this paper is organized as follows. Section 2 describes the background and related work on deep learning and parameter tuning in SE. Section 3 explains the case study problem and the

proposed tuning method investigated in this study, then Section 4 describes the experimental settings of our study, including research questions, data sets, evaluation measures and experimental design. Section 5 presents the results. Section 6 discusses implications from the results and the threats to the validity of our study. Section 7 concludes the paper and discusses the future work.

Before beginning, we digress to make two points. Firstly, just because “DE + SVM” beats deep learning in this application, this does not mean DE is always the superior method for all other software analytics applications. No learner works best over all problems [73]—the trick is to try several approaches and select the one that works best on the local data. Given the low computational cost of DE (10 minutes vs 14 hours), DEs are an obvious and low-cost candidate for exploring such alternatives.

Secondly, to enable other researchers to repeat, improve, or refute our results, all our scripts and data are freely available on-line Github¹.

2 BACKGROUND AND RELATED WORK

2.1 Why Explore Faster Software Analytics?

This section argues that avoiding slow methods for software analytics is an open and urgent issue.

Researchers and industrial practitioners now routinely make extensive use of software analytics to discover (e.g.) how long it will take to integrate the new code [17], where bugs are most likely to occur [54], who should fix the bug [2], or how long it will take to develop their code [34, 35, 50]. Large organizations like Microsoft routinely practice data-driven policy development where organizational policies are learned from an extensive analysis of large data sets collected from developers [7, 65].

But the more complex the method, the harder it is to apply the analysis. Fisher et al. [20] characterizes software analytics as a work flow that distills large quantities of low-value data down to smaller sets of higher value data. Due to the complexities and computational cost of SE analytics, “the luxuries of interactivity, direct manipulation, and fast system response are gone” [20]. They characterize modern cloud-based analytics as a throwback to the 1960s-batch processing mainframes where jobs are submitted and then analysts wait, wait, and wait for results with “little insight into what is really going on behind the scenes, how long it will take, or how much it is going to cost” [20]. Fisher et al. [20] document the issues seen by 16 industrial data scientists, one of whom remarks

“Fast iteration is key, but incompatible with the jobs are submitted and processed in the cloud. It is frustrating to wait for hours, only to realize you need a slight tweak to your feature set”.

Methods for improving the quality of modern software analytics have made this issue even more serious. There has been continuous development of new feature selection [25] and feature discovering [28] techniques for software analytics, with the most recent ones focused on deep learning methods. These are all exciting innovations with the potential to dramatically improve the quality of our software analytics tools. Yet these are all CPU/GPU-intensive methods. For instance:

- Learning control settings for learners can take days to weeks to years of CPU time [22, 64, 69].
- Lam et al. needed weeks of CPU time to combine deep learning and text mining to localize buggy files from bug reports [39].
- Gu et al. spent 240 hours of GPU time to train a deep learning based method to generate API usage sequences for given natural language query [24].

Note that the above problem is not solvable by waiting for faster CPUs/GPUs. We can no longer rely on Moore’s Law [51] to double our computational power every 18 months. Power consumption and heat dissipation issues effect block further exponential increases to CPU clock frequencies [38]. Cloud computing environments are extensively monetized so the total financial cost of training models can be prohibitive, particularly for long running tasks. For example, it would take 15 years of CPU time to learn the tuning parameters of software clone detectors proposed in [69]. Much of that CPU time can be saved if there is a faster way.

2.2 What is Deep Learning?

Deep learning is a branch of machine learning built on multiple layers of neural networks that attempt to model high level abstractions in data. According to LeCun et al. [42], deep learning methods are representation-learning methods with multiple levels of representation, obtained by composing simple but non-linear modules that each transforms the representation at one level (starting with the raw input) into a representation at a higher, slightly more abstract level. Compared to the conventional machine learning algorithms, deep learning methods are very good at exploring high-dimensional data.

By utilizing extensive computational power, deep learning has been proven to be a very powerful method by researchers in many fields [42], like computer vision and natural language processing [4, 37, 47, 60, 63]. In 2012, Convolution neural networks method won the ImageNet competition [37], which achieves half of the error rates of the best competing approaches. After that, CNN became the dominant approach for almost all recognition and detection tasks in computer vision community. CNNs are designed to process the data in the form of multiple arrays, e.g., image data. According to LeCun et al. [42], recent CNN methods are usually a huge network composed of 10 to 20 layers, hundreds of millions of weights and billions of connections between units. With advanced hardware and algorithm parallelization, training such model still need a few hours [42]. For the tasks that deal with sequential data, like text and speech, recurrent neural networks (RNNs) have been shown to work well. RNNs are found to be good at predicting the next character or word given the context. For example, Graves et al. [23] proposed to use long short-term memory (LSTM) RNNs to perform speech recognition, which achieves a test set error of 17.7% on the benchmark testing data. Sutskever et al. [63] used two multiplexlayered LSTM RNNs to translate sentences in English to French.

2.3 Deep Learning in SE

We study deep learning since, recently, it has attracted much attentions from researchers and practitioners in software community [15, 24, 39, 52, 68, 70, 71, 74, 77]. These researchers applied deep

¹<https://github.com/WeiFoo/EasyOverHard>

learning techniques to solve various problems, including defect prediction, bug localization, clone code detection, malware detection, API recommendation, effort estimation and linkable knowledge prediction.

We find that this work can be divided into two categories:

- Treat deep learning as a feature extractor, and then apply other machine learning algorithms to do further work [15, 39, 68].
- Solve problems directly with deep learning [24, 52, 70, 71, 74, 77].

2.3.1 Deep Learning as Pre-Processor. Lam et al. [39] proposed an approach to apply deep neural network in combination with rVSM to automatically locate the potential buggy files for a given bug report. By comparing it to baseline methods (Naive Bayes [32], learn-to-rank [76], BugLocator [79]), Lam et al. reported, 16.2-46.4%, 8-20.8% and 2.7-20.7% higher top-1 accuracy than baseline methods, respectively [39]. The training time for deep neural network was reported from 70 to 122 minutes for 6 projects on a computer with 32 cores 2.00GHz CPU, 126 GB memory. However, the runtime information of the baseline methods was not reported.

Wang et al. [68] applied deep belief network to automatically learn semantic features from token vectors extracted from the studied software program. After applying deep belief network to generate features from software code, Naive Bayes, ADTree and Logistic Regression methods are used to evaluate the effectiveness of feature generation, which is compared to the same learners using traditional static code features (e.g. McCabe metrics, Halstead's effort metrics and CK object-oriented code metrics [13, 26, 31, 45]). In terms of runtime, Wang et al. only report time for generating semantics features with deep belief network, which ranged from 8 seconds to 32 seconds [68]. However, the time for training and tuning deep belief network is missing. Furthermore, to compare the effectiveness of deep belief network for generating features with methods that extract traditional static code features in terms of time cost, it would be favorable to include all the time spent on feature extraction, including parsing source code, token generation and token mapping for both deep belief network and traditional methods (i.e., an end-to-end comparison).

Choetkiertikul et al. [15] proposed to apply deep learning techniques to solve effort estimation problems on user story level. Specifically, Choetkiertikul et al. [15] proposed to leverage long short-term memory (LSTM) to learn feature vectors from the title, description and comments associated with an issue report and after that, regular machine learning techniques, like CART, Random Forests, Linear Regression and Case-Based Reasoning are applied to build the effort estimation models. Experimental results show that LSTM has a significant improvement over the baseline method bag-of-words. However, no further information regarding runtime as well as experimental hardware is reported for both methods and there is no cost of this deep learning method at all.

2.3.2 Deep Learning as a Problem Solver. White et al. [70, 71] applied recurrent neural networks, a type of deep learning techniques, to address code clone detection and code suggestion. They reported, the average training time for 8 projects were ranging from 34 seconds to 2977 seconds for each epoch on a computer with two 3.3 GHz CPUs and each project required at least 30 epochs [70]. Specifically, for the *JDK* project in their experiment, it would take

25 hours on the same computer to train the models before getting prediction. For the time cost for code suggestions, authors did not mention any related information [71].

Gu et al. [24] proposed a recurrent neural network (RNN) based method, DEEPAPI, to generate API usage sequences for a given natural language query. Compared with the baseline method SWIM [57] and Lucene + UP-Miner [67], DEEPAPI improved the performance significantly. However, that improvement came at a cost: that model was trained with a Nvidia K20 GPU for 240 hours [24].

XU [74] utilized neural language model and convolution neural network (CNN) to learn word-level and document-level features to predict semantically linkable knowledge units on Stack Overflow. In terms of performance metrics, like precision, recall and F1-score, CNN method was evaluated much better than the baseline method support vector machine (SVM). However, once again, that performance improvement came at a cost: their deep learner required 14 hours to train CNN model on a 2.5GHz PC with 16 GB RAM [74].

Yuan et al. [77] proposed a deep belief network based method for malware detection on Android apps. By training and testing the deep learning model with 200 features extracted from static analysis and dynamic analysis from 500 sampled Android app, they got 96.5% accuracy for deep learning method and 80% for one baseline method, SVM [77]. However, they did not report any runtime comparison between the deep learning method and other classic machine learning methods.

Mou et al. [52] proposed a tree-based convolutional neural network for programming language processing, in which a convolution kernel is designed over programs' abstract syntax trees to capture structural information. Results show that their method achieved 94% accuracy, which is better than the baseline method RBF SVM 88.2% on program classification problem [52]. However, Mou et al. [52] did not discuss any runtime comparison between the proposed method and baseline methods.

2.3.3 Issues with Deep Learning. In summary, deep learning is used extensively in software engineering community. A common pattern in that research is to:

- Report deep learning's benefits, but not its CPU/GPU cost [15, 52, 71, 77];
- Or simply show the cost, without further analysis [24, 39, 68, 70, 74].

Since deep learning techniques cost large amount of time and computational resources to train its model, one might question whether the improvements from deep learning is worth the costs. *Are there any simple techniques that achieve similar improvements with less resource costs?* To investigate how simple methods could improve baseline methods, we select XU [74] study as a case study. The reasons are as follows:

- Most deep learning paper's baseline methods in SE are either not publicly available or too complex to implement [39, 70]. XU define their baseline methods precisely enough so others can confidently reproduce it locally. XU's baseline method is SVM learner, which is available in many machine learning toolboxes.
- Further, it is not yet common practice for deep learning researchers in SE community to share their implementations and data [15, 24, 39, 68, 70, 71], where a tiny difference may lead to a

huge difference in the results. Even though XU do not share their CNN tool, their training and testing data are available online, which can be used for our proposed method. Since the same training and testing data are used for XU's CNN and our proposed method, we can compare results of our method to their CNN results.

- Some studies do not report their runtime and experimental environment, which makes it harder for us to systematically compare our results with theirs in terms of computational costs [15, 52, 71, 77]. XU clearly report their experimental hardware and runtime, which will be easier for us compare our computational costs to theirs.

2.4 Parameter Tuning in SE

In this paper, we use DE as an optimizer to do parameter tuning for SVM, which achieves results that are competitive with deep learning. This section discusses related work on parameter tuning in SE community.

Machine learning algorithms are designed to explore the instances to learn the bias. However, most of these algorithms are controlled by parameters such as:

- The maximum allowed depth of decision tree built by CART;
- The number of trees to be built within a Random Forest.

Adjusting these parameters is called hyperparameter optimization. It is a well explored approach in other communities [9, 44]. However, in SE, such parameter optimization is not a common task (as shown in the following examples).

In the field of *defect prediction*, Fu et al. [21] surveyed hundreds of highly cited software engineering paper about defect prediction. Their observation is that most software engineering researchers did not acknowledge the impact of tunings (exceptions: [43, 64]) and use the “off-the-shelf” data miners. For example, Elish et al. [18] compared support vector machines to other data miners for the purposes of defect prediction. However, the Elish et al. paper makes no mention of any SVM tuning study [18]. More details about their survey refer to [21].

In the field of *topic modeling*, Agrawal et al. [1] investigated the impact of parameter tuning on Latent Dirichlet Allocation (LDA). LDA is a widely used technique in software engineering field to find related topics within unstructured text, like topic analytics on Stack Overflow [5] and source code analysis [10]. Agrawal et al. found that LDA suffers from conclusion instability (different input orderings can lead to very different results) that is a result of poor choice of the LDA control parameters [1]. Yet, in their survey of LDA use in SE, they found that very few researchers (4 out of 57 papers) explored the benefits of parameter tuning for LDA.

One troubling trend is that, in the few SE papers that perform tuning, they do so using methods heavily deprecated in the machine learning community. For example, two SE papers that use tuning [43, 64], apply a simple grid search to explore the potential parameter space for optimal tunings (such grid searchers run one for-loop for each parameter being optimized). However, Bergstra et al. [9] and Fu et al. [22] argue that random search methods (e.g. the differential evolution algorithm used here) are better than grid search in terms of efficiency and performance.

Table 1: Classes of knowledge unit pairs.

Class	Description
Duplicate	These two knowledge units are addressing the same question.
Direct link	One knowledge unit can help to answer the question in the other knowledge unit.
Indirect link	One knowledge provides similar information to solve the question in the other knowledge unit, but not a direct answer.
Isolated	These two knowledge units discuss unrelated questions.

3 METHOD

3.1 Research Problem

This section is an overview of the the task and methods used by XU. Their task was to predict relationships between two knowledge units (questions with answers) on Stack Overflow. Specifically, XU divided linkable knowledge unit pairs into 4 difference categories namely, *duplicate*, *direct link*, *indirect link* and *isolated*, based on its relatedness. The definition of these four categories are shown in Table 1 [74]:

In that paper, XU provided the following two methods as baselines [74]:

- TF-IDF + SVM: a multi-class SVM classifier with 36 textual features generated based on the TF and IDF values of the words in a pair of knowledge units.
- Word Embedding + SVM: a multi-class SVM classifier with word embedding generated by the word2vec model [47].

Both of these two baseline methods are compared against their proposed method, Word Embedding + CNN.

In this study, we select Word Embedding + SVM as the baseline because it uses word embedding as the input, which is the same as the Word Embedding + CNN method by XU.

3.2 Learners and Their Parameters

SVM has been proven to be a very successful method to solve text classification problem. A SVM seeks to minimize misclassification errors by selecting a boundary or hyperplane that leaves the maximum margin between positive and negative classes (where the margin is defined as the sum of the distances of the hyperplane from the closest point of the two classes [29]).

Like most machine learning algorithms, there are some parameters associated with SVM to control how it learns. In XU's experiment, they used a radial-bias function (RBF) for their SVM kernel and set γ to $1/k$, where k is 36 for TF-IDF + SVM method and 200 for Word Embedding + SVM method. For other parameters, XU mentioned that grid search was applied to optimize the SVM parameters, but no further information was disclosed.

For our work, we used the SVM module from Scikit-learn [55], a Python package for machine learning, where the parameters shown in Table. 2 are selected for tuning. Parameter C is to set the amount of regularization, which controls the tradeoff between the errors on training data and the model complexity. A small value for C will generate a simple model with more training errors, while a large value will lead to a complicated model with fewer errors.

Table 2: List of parameters tuned by this paper.

Parameters	Default	Xue et al.	Tuning Range	Description
C	1.0	unknown	[1, 50]	Penalty parameter C of the error term.
kernel	'rbf'	'rbf'	['liner', 'poly', 'rbf', 'sigmoid']	Specify the kernel type to be used in the algorithms.
gamma	1/n_features	1/200	[0, 1]	Kernel coefficient for 'rbf', 'poly' and 'sigmoid'.
coef0	0	unknown	[0, 1]	Independent term in kernel function. It is only used in 'poly' and 'sigmoid'.

Kernel is to introduce different nonlinearities into the SVM model by applying kernel functions on the input data. *Gamma* defines how far the influence of a single training example reaches, with low values meaning 'far' and high values meaning 'close'. *coef0* is an independent parameter used in sigmoid and polynomial kernel function.

As to why we used the "Tuning Range" shown in Table 2, and not some other ranges, we note that (a) those ranges include the defaults and also XU's values; (b) the results presented below show that by exploring those ranges, we achieved large gains in the performance of our baseline method. This is not to say that *larger* tuning ranges might not result in *greater* improvements. However, for the goals of this paper (to show that tuning baseline method does matter), exploring just these ranges shown in Table 2 will suffice.

3.3 Learning Word Embedding

Learning word embeddings refers to find vector representations of words such that the similarities between words can be captured by cosine similarity of corresponding vector representations. It has been shown that the words with similar semantic and syntactic are found closed to each other in the embedding space [47].

Several methods have been proposed to generate word embeddings, like skip-gram [47], GloVe [56] and PCA on the word co-occurrence matrix [41]. To replicate XU work, we used the continuous skip-gram model (word2vec), which is a unsupervised word representation learning method based on neural networks and also used by XU [74].

The skip-gram model learns vector representations of words by predicting the surrounding words in a context window. Given a sentence of words $W = w_1, w_2, \dots, w_n$, the objective of skip-gram model is to maximize the the average log probability of the surrounding words:

$$\frac{1}{n} \sum_{i=1}^n \sum_{-c \leq j \leq c, j \neq 0} \log p(w_{i+j} | w_i)$$

where c is the context window size and w_{i+j} and w_i represent surrounding words and center word, respectively. The probability of $p(w_{i+j} | w_i)$ is computed according to the softmax function:

$$p(w_o | w_I) = \frac{\exp(v_{w_o}^T v_{w_I})}{\sum_{w=1}^{|W|} \exp(v_w^T v_{w_I})}$$

where v_{w_I} and v_{w_o} are the vector representations of the input and output vectors of w , respectively. $\sum_{w=1}^{|W|} \exp(v_w^T v_{w_I})$ normalizes the inner product results across all the words. To improve the computation efficiency, Mikolove et al. [47] proposed hierarchical softmax and negative sampling techniques. More details can be found in Mikolove et al.'s study [47].

1. Given a model (e.g., SVM) with n decisions (e.g., $n = 4$), TUNER calls SAMPLE $N = 10 * n$ times. Each call generates one member of the population $pop_{i \in N}$.
2. TUNER scores each pop_i according to various objective scores o . In the case of our tuning SVM, the objective o is to maximize F1-score
3. TUNER tries to each replace pop_i with a mutant m built using Storn's differential evolution method [62]. DE extrapolates between three other members of population a, b, c . At probability p_1 , for each decision $a_k \in a$, then $m_k = a_k \vee (p_1 < \text{rand}() \wedge (b_k \vee c_k))$.
4. Each mutant m is assessed by calling EVALUATE(*model*, *prior*=*m*); i.e. by seeing what can be achieved within a goal after first assuming that *prior* = *m*.
5. To test if the mutant m is preferred to pop_i , TUNER simply compare SCORE(*m*) with SCORE(pop_i). In case of our tuning SVM, the one with higher score will be kept.
6. TUNER repeatedly loops over the population, trying to replace items with mutants, until new better mutants stop being found.
7. Return the best one in the population as the optimal tunings.

Figure 1: Procedure TUNER: strives to find "good" tunings which maximizes the objective score of the model on training and tuning data. TUNER is based on Storn's differential evolution optimizer [62].

Skip-gram's parameters control how that algorithm learns word embeddings. Those parameters include *window size* and *dimensionality of embedding space*, etc. Zuccoon et al. [80] found that embedding dimensionality and context window size have no consistent impact on retrieval model performance. However, Yang et al. [75] showed that large context window and dimension sizes are preferable to improve the performance when using CNN to solve classification tasks for Twitter. Since this work is to compare performance of tuning SVM with CNN, where skip-gram model is used to generate word vector representations for both of these methods, tuning parameter of skip-gram model is beyond the scope of this paper (but we will explore it in future work).

To train our word2vec model, 100,000 knowledge units tagged with "java" from Stack Overflow *posts* table (include titles, questions and answers) are randomly selected as a word corpus². After applying proper data processing techniques proposed by XU, like remove the unnecessary HTML tags and keep short code snippets in *code* tag, then fit the corpus into *gensim* word2vec module [58], which is a python wrapper over original word2vec package.

When converting knowledge units into vector representations, for each word w_i in the post processed knowledge unit (including title, question and answers), we query the trained word2vec model to get the corresponding word vector representation v_i . Then the

²Without further explanation, all the experiment settings, including learner algorithms, training/testing data split, etc, strictly follow XU's work.

whole knowledge unit with s words is converted to vector representation by element-wise addition, $Uv = v_1 \oplus v_2 \oplus \dots \oplus v_s$. This vector representation is used as the input data to SVM.

3.4 Tuning Algorithm

A tuning algorithm is an optimizer that drives the learner to explore the optimal parameter in a given searching space. According to our literature review, there are several searching algorithms used in SE community: *simulated annealing* [19, 46]; various *genetic algorithms* [3, 27, 30] augmented by techniques such as *differential evolution* [1, 12, 21, 22, 62], *tabu search* and *scatter search* [6, 16, 49]; *particle swarm optimization* [72]; numerous *decomposition* approaches that use heuristics to decompose the total space into small problems, then apply a *response surface methods* [36]; *NSGA-II* [78] and *NSGA-III* [48].

Of all the mentioned algorithms, the simplest are simulated annealing (SA) and differential evolution (DE), each of which can be coded in less than a page of some high-level scripting language. Our reading of the current literature is that there are more advocates for differential evolution than SA. For example, Vesterstrom and Thomsen [66] found DE to be competitive with particle swarm optimization and other GAs. DEs have already been applied before for parameter tuning in SE community to do parameter tuning (e.g. see [1, 14, 21, 22, 53]). Therefore, in this work, we adopt DE as our tuning algorithm and the main steps in DE is described in Figure 1.

4 EXPERIMENTAL SETUP

4.1 Research Questions

To systematically investigate whether tuning can improve the performance of baseline methods compared with deep learning method, we set the following three research questions:

- **RQ1:** Can we reproduce XU's baseline results (Word Embedding + SVM)?
- **RQ2:** Can DE tune a standard learner such that it outperforms XU's deep learning method?
- **RQ3:** Is tuning SVM with DE faster than XU's deep learning method?

RQ1 is to investigate whether our implementation of Word Embedding + SVM method has the similar performance with XU's baseline, which makes sure that our following analysis can be generalized to XU's conclusion. RQ2 and RQ3 lead us to investigate whether tuning SVM comparable with XU's deep learning from both performance and cost aspects.

4.2 Dataset and Experimental Design

Our experimental data comes from Stack Overflow data dump of September 2016³, where the *posts* table includes all the questions and answers posted on Stack Overflow up to date and the *postlinks* table describes the relationships between posts, e.g., *duplicate* and *linked*. As mentioned in Section 3.1, we have four different types of relationships in knowledge unit pairs. Therefore, *linked* type is further divided into *indirectly linked* and *directly linked*. Overall, four different types of data are generated according the following rules [74]:

³<https://archive.org/details/stackexchange>

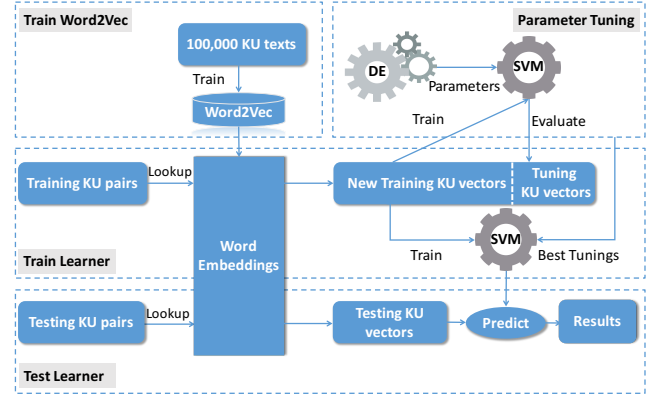


Figure 2: The overall workflow of building knowledge units predictor with tuned SVM

- Randomly select a pair of posts from the *postlinks* table, if the value in *PostLinkTypeId* field for this pair of posts is 3, then this pair of posts is *duplicate* posts. Otherwise they're *directly linked* posts.
- Randomly select a pair of posts from the *posts* table, if this pair of posts is linkable from each other according to *postlinks* table and the distance between them are greater than 2 (which means they are not duplicate or directly linked posts), then this pair of posts is *indirectly linked*. If they're not linkable, then this pair of posts is *isolated*.

In this work, we use the same training and testing knowledge unit pairs as XU [74]⁴, where 6,400 pairs of knowledge units for training and 1,600 pairs for testing. And each type of linked knowledge units accounts for 1/4 in both training and testing data. The reasons that we used the same training and testing data as XU are:

- It is to ensure that performance of our baseline method is as closed to XU's as possible.
- Since deep learning method is way complicated compared to SVM and a little difference in implementations might lead to different results. To fairly compare with XU's result, we can use the performance scores of CNN method from XU's study [74] without any implementation bias introduced.

For training word2vec model, we randomly select 100,000 knowledge units (title, question body and all the answers) from *posts* table that are related to "java". After that, all the training/tuning/testing knowledge units used in this paper are converted into word embedding representations by looking up each word in word2vec model as described in Section 3.3.

As seen in Figure 2, instead of using all the 6,400 knowledge units as training data, we split the original training data into new training data and tuning data, which are used during parameter tuning procedure for training SVM and evaluating candidate parameters offered by DE. Afterwards, the new training data is again fitted into the SVM with the optimal parameters found by DE and finally the performance of the tuned SVM will be evaluated on the testing data.

⁴<https://github.com/XBWer/ASEDataset>

To reduce the potential variance caused by how the original training data is divided, *10-fold cross-validation* is performed. Specifically, each time one fold with 640 knowledge units pairs is used as the tuning data, and the remaining folds with 5760 knowledge units are used as the new training data, then the output SVM model will be evaluated on the testing data. Therefore, all the performance scores reported below are averaged values over 10 runs.

In this study, we use Wilcoxon single ranked test to statistically compare the differences between tuned SVM and untuned SVM. Specifically, the Benjamini-Hochberg (BH) adjusted p-value is used to test whether a difference is statistically significant at the level of 0.05 [8]. To measure the effect size of performance scores between tuned SVM and untuned SVM, we compute Cliff's δ that is a non-parametric effect size measure [59]. As Romano et al. suggested, we evaluate the magnitude of the effect size as follows: negligible ($|\delta| < 0.147$), small ($0.147 < |\delta| < 0.33$), medium ($0.33 < |\delta| < 0.474$), and large ($0.474 \leq |\delta|$) [59].

4.3 Evaluation Metrics

When evaluating the performance of tuning SVM on the multi-class linkable knowledge units prediction problem, consistent with XU [74], we use accuracy, precision, recall and F1-score as the evaluation metrics.

Table 3: Confusion matrix.

		Classified as			
		C_1	C_2	C_3	C_4
Actual	C_1	c_{11}	c_{12}	c_{13}	c_{14}
	C_2	c_{21}	c_{22}	c_{23}	c_{24}
	C_3	c_{31}	c_{32}	c_{33}	c_{34}
	C_4	c_{41}	c_{42}	c_{43}	c_{44}

Given a multi-classification problem with true labels C_1, C_2, C_3 and C_4 , we can generate a confusion matrix like Table 3, where the value of c_{ii} represents the number of instances that are correctly classified by the learner for class C_i .

Accuracy is defined as the number of correctly classified knowledge units over the total number of knowledge units, i.e.,

$$accuracy = \frac{\sum_i c_{ii}}{\sum_i \sum_j c_{ij}}$$

where $\sum_i \sum_j c_{ij}$ is the total number of knowledge units. For a given type of knowledge units, C_j , the precision is defined as probability of knowledge units pairs correctly classified as C_j over the number of knowledge unit pairs classified as C_j and recall is defined as the percentage of all C_j knowledge unit pairs correctly classified. F1-score is the harmonic mean of recall and precision. Mathematically, precision, recall and F1-score of the learner for class C_j can be denoted as follows:

$$\begin{aligned} prec_j &= precision_j = \frac{c_{jj}}{\sum_i c_{ij}} \\ pd_j &= recall_j = \frac{c_{jj}}{\sum_i c_{ji}} \\ F1_j &= 2 * pd_j * prec_j / (pd_j + prec_j) \end{aligned}$$

Where $\sum_i c_{ij}$ is the predicted number of knowledge units in class C_j and $\sum_i c_{ji}$ is the actual number of knowledge units in class C_j .

Recall from Algorithm 1 that we call differential evolution once for each optimization goal. Generally, this goal depends on which

metric is most important for the business case. In this work, we use F1 to score the candidate parameters because it controls the trade-off between precision and recall, which is also consistent with XU [74] and is also widely used in software engineering community to evaluate classification results [21, 33, 46, 68].

5 RESULTS

In this section, we present our experimental results. To answer research questions raised in Section 4.1, we conducted two experiments:




Figure 3: Score delta between our SVM with XU's SVM in [74] in terms of precision, recall and F1-score. Positive values mean our SVM is better than XU's SVM in terms of different measures; Otherwise, XU's SVM is better.

- Compare performance of Word Embedding + SVM method in XU [74] and our implementation;
- Compare performance of our tuning SVM with DE method with XU's CNN deep learning method.

Since we used the same training and testing data sets provided by XU [74] and conducted our experiment in the same procedure and evaluated methods using the performance measures, we simply used the results reported in the work by XU [74] for the performance comparison.

RQ1: Can we reproduce XU's baseline results (Word Embedding + SVM)?

This first question is important to our work since, without the original tool released by XU, we need to insure that our reimplementation of their baseline method (WordEmbedding + SVM) has a similar performance to their work. Accordingly, we carefully follow XU's procedure [74]. We use the SVM learner from scikit-learn with the setting $\gamma = \frac{1}{200}$ and *kernel* = "rbf", which are used by XU. After that, the same training and testing knowledge unit pairs are applied to SVM.

Table 4 and Figure 3 show the performance scores and corresponding score delta between our implementation of WordEmbedding + SVM with XU's in terms of accuracy⁵, precision, recall and F1-score. As we can see, when predicting these four different types of relatedness between knowledge unit pairs, our Word Embedding

⁵XU just report overall accuracy, not for each class, hence it is missing in this table.




Figure 4: Score delta between tuned SVM and CNN method [74] in terms of precision, recall and F1-score. Positive values mean tuned SVM is better than CNN in terms of different measures; Otherwise, CNN is better.

Table 4: Comparison of our baseline method with XU's. The best scores are marked in bold.

Metrics	Methods	Duplicate	Direct Link	Indirect Link	Isolated	Overall
Precision	Our SVM	0.724	0.514	0.779	0.601	0.655
	XU's SVM	0.611	0.560	0.787	0.676	0.659
Recall	Our SVM	0.525	0.492	0.970	0.645	0.658
	XU's SVM	0.725	0.433	0.980	0.538	0.669
F1-score	Our SVM	0.609	0.503	0.864	0.622	0.650
	XU's SVM	0.663	0.488	0.873	0.600	0.656
Accuracy	Our SVM	0.525	0.493	0.970	0.645	0.658
	XU's SVM	-	-	-	-	0.669

+ SVM method has very similar performance scores to the baseline method reported by XU in [74], with the maximum difference less than 0.2. Except for Duplicate class, where our baseline has a higher precision (i.e., 0.724 v.s. 0.611) but a lower recall (i.e., 0.525 v.s. 0.725).

Figure 3 presents the same results in a graphical format. Any bar above zero means that our implementation has a better performance score than XU's on predicting that specific knowledge unit relatedness class. As we can see, most of the differences ($\frac{8}{12}$) are within 0.05 and the score delta of overall performance shows that our implementation is a little worse than XU's implementation. For this chart we conclude that:

Overall, our reimplement of WordEmbedding + SVM has very similar performance in all the evaluated metrics compared to the baseline method reported in XU's study [74].

The significance of this conclusion is that, moving forward, we are confident that we can use our reimplement of WordEmbedding+SVM as a valid surrogate for the baseline method of XU.

RQ2: Can DE tune a standard learner such that it outperforms XU's deep learning method?

To answer this question, we run the workflow of Figure 2, where DE is applied to find the optimal parameters of SVM based on the training and tuning data. The optimal tunings are then applied on




Figure 5: Score delta between the tuned SVM and XU's baseline SVM in terms of precision, recall and F1-score. Positive values mean the tuned SVM is better than XU's SVM in terms of different measures; Otherwise, XU's SVM is better.

Table 5: Comparison of the tuned SVM with XU's CNN method. The best scores are marked in bold.

Metrics	Methods	Duplicate	Direct Link	Indirect Link	Isolated	Overall
Precision	XU's SVM	0.611	0.560	0.787	0.676	0.658
	XU's CNN	0.898	0.758	0.840	0.890	0.847
	Tuned SVM	0.885	0.851	0.944	0.903	0.896
Recall	XU's SVM	0.725	0.433	0.980	0.538	0.669
	XU's CNN	0.898	0.903	0.773	0.793	0.842
	Tuned SVM	0.860	0.828	0.995	0.905	0.897
F1-score	XU's SVM	0.663	0.488	0.873	0.600	0.656
	XU's CNN	0.898	0.824	0.805	0.849	0.841
	Tuned SVM	0.878	0.841	0.969	0.909	0.899

the SVM model and the built learner is evaluated on testing data. Note that, in this study, since we mainly focus on precision, recall and F1-score measures where F1-score is the harmonic mean of precision and recall, we use F1-score as the tuning goal for DE. In other words, when tuning parameters, DE expects to find a pair of candidate parameters that maximize F1-score.

Table 5 presents the performance scores of XU's baseline, XU's CNN method and Tuned SVM for all metrics. The highest score for each relatedness class is marked in bold. Note that: Without tuning, XU's CNN method outperforms the baseline SVM in $\frac{10}{12}$ evaluation metrics across all four classes. The largest performance improvement is 0.47 for recall on Direct Link class. Note that this result is consistent with XU's conclusion that their CNN method is superior to standard SVM. After tuning SVM, the deep learning method has no such advantage. Specifically, CNN has advantage over tuned SVM in $\frac{4}{12}$ evaluation metrics across all four classes. Even when CNN performs better than our tuning SVM method, the largest difference is 0.065 for Recall on Direct_Link class, which is less than 0.1.

Figure 4 presents the same results in a graphical format. Any bar above zero indicates that tuned SVM has a better performance score than CNN. In this figure: CNN has a slightly better performance on Duplicate class for precision, recall and F1-score and a higher

recall on Direct link class. Across all of Figure 4, in $\frac{8}{12}$ evaluation scores, Tuned SVM has better performance scores than CNN, with the largest delta of 0.222.

Figure 5 compares the performance delta of tuned SVM with XU's untuned SVM. We note that DE-based parameter tuning never degrades SVM's performance (since there are no negative values in that chart). Tuning dramatically improves scores on predicting some classes of KU relatedness. For example, the recall of predicting Direct_link is increased from 0.433 to 0.903, which is 108% improvement over XU's untuned SVM (To be fair for XU, it is still 84% improvement over our untuned SVM). At the same time, the corresponding precision and F1 scores of predicting Direct_Link are increased from 0.560 to 0.851 and 0.488 to 0.841, which are 52% and 72% improvement over XU's original report[74], respectively. A similar pattern can also be observed in Isolated class. On average, tuning helps improve the performance of XU's SVM by 0.238, 0.228 and 0.227 in terms of precision, recall and F1-score for all four KU relatedness classes. Figure 6 compares the tuned SVM with our untuned SVM. We note that we get the similar patterns that observed in Figure 5. All the bars are above zero, etc.

Based on the performance scores in Table 5 and score delta in Figure 4, Figure 5 and Figure 6, we can see that:

- Parameter tuning can dramatically improve the performance of Word Embedding + SVM (the baseline method) for the multi-class KU relatedness prediction task;
- With the optimal tunings, the traditional machine learning method, SVM, if not better, is at least comparable with deep learning methods (CNN).

When discussing this result with colleagues, we are sometimes asked for a statistical analysis that confirms the above finding. However, due the lack of evaluation score distributions of the CNN method in [74], we cannot compare their single value with our results from 10 repeated runs. However, according to Wilcoxon signed rank test over 10 runs results, tuned SVM performs statistically better than our untuned SVM in terms of all evaluation measures on all four classes ($p < 0.05$). According to Cliff δ values, the magnitude of difference between tuned SVM and our untuned SVM is not trivial ($|\delta| > 0.147$) for all evaluation measures.

Overall, the experimental results and our analysis indicate that:

In the evaluation conducted here, the deep learning method, CNN, does not have any performance advantage over our tuning approach.

RQ3: Is tuning SVM with DE faster than XU's deep learning method?

When comparing the runtime of two learning methods, it obviously should be conducted under the same hardware settings. Since we adopt the CNN evaluation scores from [74], we can not run on our tuning SVM experiment under the exactly same system settings. To allow readers to have a objective comparison, we provide the experimental environment as shown in Table 6. To obtain the runtime of tuning SVM, we recorded the start time and end time of the program execution, including parameter tuning, training model and testing model.

According to XU, it took 14 hours to train their CNN model into a low loss convergence ($< e^{-3}$) [74]. Our work, on the other hand




Figure 6: Score delta between tuned SVM and Our Untuned SVM in terms of precision, recall and F1-score. Positive values mean the tuned SVM is better than our untuned SVM in terms of different measures; Otherwise, our SVM is better.

Table 6: Comparison of experimental environment

Methods	OS	CPU	RAM
Tuning SVM	MacOS 10.12	Intel Core i5 2.7 GHz	8 GB
CNN	Windows 7	Intel Core i7 2.5 GHz	16 GB

only takes 10 minutes to run SVM with parameter tuning by DE on a similar environment. That is, the simple parameter tuning method on SVM is 84X faster than XU's deep learning method.

Compared to CNN method, tuning SVM is about 84X faster in terms of model building.

The significance of this finding is that, in this case study, CNN was neither better in performance scores (see RQ2) nor runtimes. CNN's extra runtimes are a particular concern since (a) they are very long; and (b) these would be incurred anytime researchers wants to update the CNN model with new data or wanted to validate the XU result.

6 DISCUSSION

6.1 Why DE+SVM works?

Parameter tuning matters. As mentioned in Section 2.4, the default parameter values set by the algorithm designers could generate a good performance on average but may not guarantee the best performance for the local data [9, 21]. Given that, it is most strange to report that most SE researchers ignore the impacts of parameter tuning when they utilize various machine learning methods to conduct software analytic (evidence: see our reviews in [1, 21, 22]). The conclusion of this work must be to stress the importance of this kind of tuning, using local data, for any future software analytics study.

Better explore the searching space. It turns out that one exception to our statement that "most researchers do not tune" is the XU study. In that work, they unsuccessfully perform parameter tuning, but with with grid search. In such a grid search, for N parameters to be tuned, N for loops are created to run over a

range of settings for each parameter. While a widely used method, it is often deprecated. For example, Bergstra et al.[9] note that grid search jumps through different parameter settings between some *min* and *max* values of pre-defined tuning range. They warn that such jumps may actually skip over the critical tuning values. On the other hand, DE tuning values are adjusted based on better candidates from previous generations. Hence DE is more likely than grid search to “fill in the gaps” between the initialized values.

That said, although DE +SVM works in this study, it does not mean DE is the best parameter tuner for all SE tasks. We encourage more researchers to explore faster and more effective parameter tuners in this direction.

6.2 Implication

Beyond the specifics of this case study, what general principles can we take from the above work?

Understand the task. One reason to try different tools for the same task is to better understand the task. The more we understand a task, the better we can match tools to that task. Tools that are poorly matched to task are usually complex and/or slow to execute. In the case study of this paper, we would say that

- Deep learning is a poor match to the task of predicting whether two questions posted on Stack Overflow are semantically linkable since it is so slow;
- Differential evolution tuning SVM is a much better match since it is so fast and obtain competitive performance.

That said, it is important to stress that the point of this study is not to deprecate deep learning. There are many scenarios where we believe deep learning would be a natural choice (e.g. when analyzing complex speech or visual data). In SE, it is still an open research question that in which scenario deep learning is the best choice. Results from this paper show that, at least for classification tasks like knowledge unit relatedness classification on Stack Overflow, deep learning does not have much advantage over well tuned conventional machine learning methods. However, as we better understand SE tasks, deep learning could be used to address more SE problems, which require more advanced artificial intelligence.

Treat resource constraints as design challenges. As a general engineering principle, we think it insightful to consider the resource cost of a tool before applying it. It turns out that this is a design pattern used in contemporary industry. According to Calero and Pattini [11], many current commercial redesigns are motivated (at least in part) by arguments based on sustainability (i.e. using fewer resources to achieve results). In fact, they say that managers used sustainability-based redesigns to motivate extensive cost-cutting opportunities.

6.3 Threats to Validity

Threats to **internal validity** concern the consistency of the results obtained from the result. In our study, to investigate how tuning can improve the performance of baseline methods and how well it perform compared with deep learning method. We select XU’s Word Embedding + SVM baseline method as a case study. Since the original implementation of Word Embedding + SVM (baseline 2 method in [74]) is not publicly available, we have to reimplement our version of Word Embedding + SVM as the baseline method in

this study. As shown in RQ1, our implementation similar results to XU’s on the same data sets. Hence, we believe that our implementation reflect the original baseline method in Xu’s study [74].

Threats to **external validity** represent if the results are of relevance for other cases, or the ability to generalize the observations in a study. In this study, we compare our tuning baseline method with deep learning method, CNN, in terms of precision, recall, F1-score and accuracy. The experimental results are quite consistent for this knowledge units relatedness prediction task. Nonetheless, we do not claim that our findings can be generalized to all software analytics tasks. However, those other software analytics tasks often apply deep learning methods on classification tasks [15, 68] and so it is quite possible that the methods of this paper (i.e., DE-based parameter tuning) would be widely applicable, elsewhere.

7 CONCLUSION

In this paper, we perform a comparative study to investigate how tuning can improve the baseline method compared with state-of-the-art deep learning method for predicting knowledge units relatedness on Stack Overflow. Our experimental results show that:

- Tuning improves the performance of baseline methods. At least for Word Embedding + SVM (baseline in [74]) method, if not better, it performs as well as the proposed CNN method in [74].
- The baseline method with parameter tuning runs much faster than complicated deep learning. In this study, tuning SVM runs 84X faster than CNN method.

8 ADDENDUM

As this paper was going to press we learned of a new deep learning methods that, according to its creators, runs 20 times faster than standard deep learning [61]. Note that in that paper, the authors say their faster method does not produce better results— in fact, their method generated solutions that were a small fraction worse than “classic” deep learning. Hence, that paper does not invalidate our result since (a) our DE-based method sometimes produced *better* results than classic deep learning and (b) our DE runs 84 times faster (i.e. much faster than those reported in [61]).

That said, this new fast deep learning method deserves our close attention since, using it, we conjecture that our DE tools could solve an open problem in the deep learning community; i.e. how to find the best configurations inside a deep learner faster.

Based on the results of this study, we recommend that before applying deep learning method on SE tasks, implement simpler techniques. These simpler methods could be used, at the very least, for comparisons against a baseline. In this particular case of deep learning vs DE, the extra computational effort is so very minor (10 minutes on top of 14 hours), that such a “try-with-simpler” should be standard practice.

As to the future work, we will explore more simple techniques to solve SE tasks and also investigate how deep learning techniques could be applied effectively in software engineering field.

ACKNOWLEDGEMENTS

The work is partially funded by an NSF award #1302169.

REFERENCES

- [1] Amritanshu Agrawal, Wei Fu, and Tim Menzies. 2016. What is wrong with topic modeling?(and how to fix it using search-based se). *arXiv preprint arXiv:1608.08176* (2016).
- [2] John Anvik, Lyndon Hiew, and Gail C Murphy. 2006. Who should fix this bug?. In *Proceedings of the 28th International Conference on Software Engineering*. ACM, 361–370.
- [3] Andrea Arcuri and Gordon Fraser. 2011. On parameter tuning in search based software engineering. In *International Symposium on Search Based Software Engineering*. Springer, 33–47.
- [4] Itamar Arel, Derek C Rose, and Thomas P Karnowski. 2010. Research Frontier: Deep Machine Learning—a New Frontier in Artificial Intelligence Research. *IEEE Computational Intelligence Magazine* 5, 4 (2010), 13–18.
- [5] Anton Barua, Stephen W Thomas, and Ahmed E Hassan. 2014. What are developers talking about? an analysis of topics and trends in stack overflow. *Empirical Software Engineering* 19, 3 (2014), 619–654.
- [6] Ricardo P Beausoleil. 2006. “MOSS” multiobjective scatter search applied to non-linear multiple criteria optimization. *European Journal of Operational Research* 169, 2 (2006), 426–449.
- [7] Andrew Begel and Thomas Zimmermann. 2014. Analyze this! 145 questions for data scientists in software engineering. In *Proceedings of the 36th International Conference on Software Engineering*. ACM, 12–23.
- [8] Yoav Benjamini and Yosef Hochberg. 1995. Controlling the false discovery rate: a practical and powerful approach to multiple testing. *Journal of the royal statistical society. Series B (Methodological)* (1995), 289–300.
- [9] James Bergstra and Yoshua Bengio. 2012. Random search for hyper-parameter optimization. *Journal of Machine Learning Research* 13, Feb (2012), 281–305.
- [10] David Binkley, Daniel Heinz, Dawn Lawrie, and Justin Overfelt. 2014. Understanding LDA in source code analysis. In *Proceedings of the 22nd International Conference on Program Comprehension*. ACM, 26–36.
- [11] Coral Calero and Mario Piattini. 2015. Green in Software Engineering. (2015).
- [12] José M Chaves-González and Miguel A Pérez-Toledano. 2015. Differential evolution with Pareto tournament for the multi-objective next release problem. *Appl. Math. Comput.* 252 (2015), 1–13.
- [13] Shyam R Chidamber and Chris F Kemerer. 1994. A metrics suite for object oriented design. *IEEE Transactions on Software Engineering* 20, 6 (1994), 476–493.
- [14] Ibtissem Chiha, J Ghabri, and Noureddine Liouane. 2012. Tuning PID controller with multi-objective differential evolution. In *2012 5th International Symposium on Communications, Control and Signal Processing*. IEEE, 1–4.
- [15] Morakot Choetkiertikul, Hoa Khanh Dam, Truyen Tran, Trang Pham, Aditya Ghose, and Tim Menzies. 2016. A deep learning model for estimating story points. *arXiv preprint arXiv:1609.00489* (2016).
- [16] Anna Corazza, Sergio Di Martino, Filomena Ferrucci, Carmine Gravino, Federica Sarro, and Emilia Mendes. 2013. Using tabu search to configure support vector regression for effort estimation. *Empirical Software Engineering* 18, 3 (2013), 506–546.
- [17] Jacek Czerwinka, Rajiv Das, Nachiappan Nagappan, Alex Tarvo, and Alex Teterev. 2011. Crane: Failure prediction, change analysis and test prioritization in practice—experiences from windows. In *Proceedings of the 4th IEEE International Conference on Software Testing, Verification and Validation*. IEEE, 357–366.
- [18] Karim O Elish and Mahmoud O Elish. 2008. Predicting defect-prone software modules using support vector machines. *Journal of Systems and Software* 81, 5 (2008), 649–660.
- [19] Martin S Feather and Tim Menzies. 2002. Converging on the optimal attainment of requirements. In *Proceedings of the 10th Anniversary IEEE Joint International Conference on Requirements Engineering*. IEEE, 263–270.
- [20] Danyel Fisher, Rob DeLine, Mary Czerwinski, and Steven Drucker. 2012. Interactions with big data analytics. *interactions* 19, 3 (2012), 50–59.
- [21] Wei Fu, Tim Menzies, and Xipeng Shen. 2016. Tuning for software analytics: Is it really necessary? *Information and Software Technology* 76 (2016), 135–146.
- [22] Wei Fu, Vivek Nair, and Tim Menzies. 2016. Why is differential evolution better than grid search for tuning defect predictors? *arXiv preprint arXiv:1609.02613* (2016).
- [23] Alex Graves, Abdel-rahman Mohamed, and Geoffrey Hinton. 2013. Speech recognition with deep recurrent neural networks. In *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*. IEEE, 6645–6649.
- [24] Xiaodong Gu, Hongyu Zhang, Dongmei Zhang, and Sunghun Kim. 2016. Deep API learning. In *Proceedings of the 2016 24th ACM SIGSOFT International Symposium on Foundations of Software Engineering*. ACM, 631–642.
- [25] Mark A Hall and Geoffrey Holmes. 2003. Benchmarking attribute selection techniques for discrete class data mining. *IEEE Transactions on Knowledge and Data Engineering* 15, 6 (2003), 1437–1447.
- [26] Maurice Howard Halstead. 1977. *Elements of software science*. Vol. 7. Elsevier New York.
- [27] Mark Harman. 2007. The current state and future of search based software engineering. In *2007 Future of Software Engineering*. IEEE Computer Society, 342–357.
- [28] Tian Jiang, Lin Tan, and Sunghun Kim. 2013. Personalized defect prediction. In *Proceedings of the 28th IEEE/ACM International Conference on Automated Software Engineering*. IEEE, 279–289.
- [29] Thorsten Joachims. 1998. Text categorization with support vector machines: Learning with many relevant features. In *European Conference on Machine Learning*. Springer, 137–142.
- [30] Bryan F Jones, H-H Sthamer, and David E Eyres. 1996. Automatic structural testing using genetic algorithms. *Software Engineering Journal* 11, 5 (1996), 299–306.
- [31] Dennis Kafura and Geeredy R Reddy. 1987. The use of software complexity metrics in software maintenance. *IEEE Transactions on Software Engineering* 3 (1987), 335–343.
- [32] Dongsun Kim, Yida Tao, Sunghun Kim, and Andreas Zeller. 2013. Where should we fix this bug? a two-phase recommendation model. *IEEE Transactions on Software Engineering* 39, 11 (2013), 1597–1610.
- [33] Sunghun Kim, E James Whitehead Jr, and Yi Zhang. 2008. Classifying software changes: Clean or buggy? *IEEE Transactions on Software Engineering* 34, 2 (2008), 181–196.
- [34] Ekrem Kocaguneli, Tim Menzies, Ayse Bener, and Jacky W Keung. 2012. Exploiting the essential assumptions of analogy-based effort estimation. *IEEE Transactions on Software Engineering* 38, 2 (2012), 425–438.
- [35] Ekrem Kocaguneli, Tim Menzies, and Jacky W Keung. 2012. On the value of ensemble effort estimation. *IEEE Transactions on Software Engineering* 38, 6 (2012), 1403–1416.
- [36] Joseph Krall, Tim Menzies, and Misty Davies. 2015. Gale: Geometric active learning for search-based software engineering. *IEEE Transactions on Software Engineering* 41, 10 (2015), 1001–1018.
- [37] Alex Krizhevsky, Ilya Sutskever, and Geoffrey E Hinton. 2012. Imagenet classification with deep convolutional neural networks. In *Advances in Neural Information Processing Systems*. 1097–1105.
- [38] Rakesh Kumar, Keith I Farkas, Norman P Jouppi, Parthasarathy Ranganathan, and Dean M Tullsen. 2003. Single-ISA heterogeneous multi-core architectures: The potential for processor power reduction. In *Proceedings of the 36th Annual IEEE/ACM International Symposium on Microarchitecture*. IEEE, 81–92.
- [39] An Ngoc Lam, Anh Tuan Nguyen, Hoan Anh Nguyen, and Tien N Nguyen. 2015. Combining deep learning with information retrieval to localize buggy files for bug reports (n). In *Proceedings of the 2015 30th IEEE/ACM International Conference on Automated Software Engineering*. IEEE, 476–481.
- [40] Quoc V Le. 2013. Building high-level features using large scale unsupervised learning. In *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*. IEEE, 8595–8598.
- [41] Rémi Lebrete and Ronan Collobert. 2013. Word emdeddings through hellinger PCA. *arXiv preprint arXiv:1312.5542* (2013).
- [42] Yann LeCun, Yoshua Bengio, and Geoffrey Hinton. 2015. Deep learning. *Nature* 521, 7553 (2015), 436–444.
- [43] Stefan Lessmann, Bart Baesens, Christophe Mues, and Swantje Pietsch. 2008. Benchmarking classification models for software defect prediction: A proposed framework and novel findings. *IEEE Transactions on Software Engineering* 34, 4 (2008), 485–496.
- [44] Lisha Li, Kevin Jamieson, Giulia DeSalvo, Afshin Rostamizadeh, and Ameet Talwalkar. 2016. Hyperband: a novel bandit-based approach to hyperparameter optimization. *arXiv preprint arXiv:1603.06560* (2016).
- [45] Thomas J McCabe. 1976. A complexity measure. *IEEE Transactions on Software Engineering* 4 (1976), 308–320.
- [46] Tim Menzies, Jeremy Greenwald, and Art Frank. 2007. Data mining static code attributes to learn defect predictors. *IEEE Transactions on Software Engineering* 33, 1 (2007).
- [47] Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. 2013. Distributed representations of words and phrases and their compositionality. In *Advances in Neural Information Processing Systems*. 3111–3119.
- [48] Mohamed Wiem Mkaouer, Marouane Kessentini, Slim Bechikh, Kalyanmoy Deb, and Mel Ó Cinnéide. 2014. High dimensional search-based software engineering: finding tradeoffs among 15 objectives for automating software refactoring using NSGA-III. In *Proceedings of the 2014 Annual Conference on Genetic and Evolutionary Computation*. ACM, 1263–1270.
- [49] Julian Molina, Manuel Laguna, Rafael Marti, and Rafael Caballero. 2007. SSPMO: A scatter tabu search procedure for non-linear multiobjective optimization. *INFORMS Journal on Computing* 19, 1 (2007), 91–100.
- [50] Kjetil Molokken and Magen Jorgensen. 2003. A review of software surveys on software effort estimation. In *Proceedings of the 2003 International Symposium on Empirical Software Engineering*. IEEE, 223–230.
- [51] Gordon E Moore and others. 1998. Cramming more components onto integrated circuits. *Proc. IEEE* 86, 1 (1998), 82–85.
- [52] Lili Mou, Ge Li, Lu Zhang, Tao Wang, and Zhi Jin. 2016. Convolutional Neural Networks over Tree Structures for Programming Language Processing. In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*. AAAI Press, 1287–1293.

- [53] Mahamed GH Omran, Andries Petrus Engelbrecht, and Ayed Salman. 2005. Differential evolution methods for unsupervised image classification. In *2005 IEEE Congress on Evolutionary Computation*, Vol. 2. IEEE, 966–973.
- [54] Thomas J Ostrand, Elaine J Weyuker, and Robert M Bell. 2004. Where the bugs are. In *ACM SIGSOFT Software Engineering Notes*, Vol. 29. ACM, 86–96.
- [55] Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, and others. 2011. Scikit-learn: machine learning in Python. *Journal of Machine Learning Research* 12, Oct (2011), 2825–2830.
- [56] Jeffrey Pennington, Richard Socher, and Christopher D Manning. 2014. Glove: global vectors for word representation. In *EMNLP*, Vol. 14. 1532–1543.
- [57] Mukund Raghothaman, Yi Wei, and Youssef Hamadi. 2016. SWIM: synthesizing what I mean: code search and idiomatic snippet synthesis. In *Proceedings of the 38th International Conference on Software Engineering*. ACM, 357–367.
- [58] Radim Rehurek and Petr Sojka. 2010. Software framework for topic modelling with large corpora. In *In Proceedings of the LREC 2010 Workshop on New Challenges for NLP Frameworks*. Citeseer.
- [59] Jeanine Romano, Jeffrey D Kromrey, Jesse Coraggio, Jeff Skowronek, and Linda Devine. 2006. Exploring methods for evaluating group differences on the NSSE and other surveys: Are the t-test and Cohen's d indices the most appropriate choices. In *Annual Meeting of the Southern Association for Institutional Research*.
- [60] Jürgen Schmidhuber. 2015. Deep learning in neural networks: An overview. *Neural networks* 61 (2015), 85–117.
- [61] Ryan Spring and Anshumali Shrivastava. 2017. Scalable and sustainable deep learning via randomized hashing. In *Proceedings of the 23rd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*. ACM.
- [62] R. Storn and K. Price. 1997. Differential evolution—a simple and efficient heuristic for global optimization over continuous spaces. *Journal of global optimization* 11, 4 (1997), 341–359.
- [63] Ilya Sutskever, Oriol Vinyals, and Quoc V Le. 2014. Sequence to sequence learning with neural networks. In *Advances in Neural Information Processing Systems*. 3104–3112.
- [64] Chakkrit Tantithamthavorn, Shane McIntosh, Ahmed E Hassan, and Kenichi Matsumoto. 2016. Automated parameter optimization of classification techniques for defect prediction models. In *Proceedings of the 38th International Conference on Software Engineering*. ACM, 321–332.
- [65] Christopher Theisen, Kim Herzig, Patrick Morrison, Brendan Murphy, and Laurie Williams. 2015. Approximating attack surfaces with stack traces. In *Proceedings of the 37th International Conference on Software Engineering—Volume 2*. IEEE, 199–208.
- [66] Jakob Vesterstrom and Rene Thomsen. 2004. A comparative study of differential evolution, particle swarm optimization, and evolutionary algorithms on numerical benchmark problems. In *Proceedings of the 2004 Congress on Evolutionary Computation*, Vol. 2. IEEE, 1980–1987.
- [67] Jue Wang, Yingnong Dang, Hongyu Zhang, Kai Chen, Tao Xie, and Dongmei Zhang. 2013. Mining succinct and high-coverage API usage patterns from source code. In *Proceedings of the 10th Working Conference on Mining Software Repositories*. IEEE, 319–328.
- [68] Song Wang, Taiyue Liu, and Lin Tan. 2016. Automatically learning semantic features for defect prediction. In *Proceedings of the 38th International Conference on Software Engineering*. ACM, 297–308.
- [69] Tiantian Wang, Mark Harman, Yue Jia, and Jens Krinke. 2013. Searching for better configurations: a rigorous approach to clone evaluation. In *Proceedings of the 2013 9th Joint Meeting on Foundations of Software Engineering*. ACM, 455–465.
- [70] Martin White, Michele Tufano, Christopher Vendome, and Denys Poshyvanyk. 2016. Deep learning code fragments for code clone detection. In *Proceedings of the 31st IEEE/ACM International Conference on Automated Software Engineering*. ACM, 87–98.
- [71] Martin White, Christopher Vendome, Mario Linares-Vásquez, and Denys Poshyvanyk. 2015. Toward deep learning software repositories. In *Proceedings of the 12th Working Conference on Mining Software Repositories*. IEEE, 334–345.
- [72] Andreas Windisch, Stefan Wappler, and Joachim Wegener. 2007. Applying particle swarm optimization to software testing. In *Proceedings of the 9th Annual Conference on Genetic and Evolutionary Computation*. ACM, 1121–1128.
- [73] David H Wolpert. 1996. The lack of a priori distinctions between learning algorithms. *Neural computation* 8, 7 (1996), 1341–1390.
- [74] Bowen Xu, Deheng Ye, Zhenchang Xing, Xin Xia, Guibin Chen, and Shanping Li. 2016. Predicting semantically linkable knowledge in developer online forums via convolutional neural network. In *Proceedings of the 31st IEEE/ACM International Conference on Automated Software Engineering*. ACM, 51–62.
- [75] Xiao Yang, Craig Macdonald, and Iadh Ounis. 2016. Using word embeddings in twitter election classification. *arXiv preprint arXiv:1606.07006* (2016).
- [76] Xin Ye, Razvan Bunescu, and Chang Liu. 2014. Learning to rank relevant files for bug reports using domain knowledge. In *Proceedings of the 22nd ACM SIGSOFT International Symposium on Foundations of Software Engineering*. ACM, 689–699.
- [77] Zhenlong Yuan, Yongqiang Lu, Zhaoguo Wang, and Yibo Xue. 2014. Droidsec: deep learning in android malware detection. In *ACM SIGCOMM Computer Communication Review*, Vol. 44. ACM, 371–372.
- [78] Yuanyuan Zhang, Mark Harman, and S. Afshin Mansouri. 2007. The multi-objective next release problem. In *Proceedings of the 9th Annual Conference on Genetic and Evolutionary Computation*. 1129–1137.
- [79] Jian Zhou, Hongyu Zhang, and David Lo. 2012. Where should the bugs be fixed?—more accurate information retrieval-based bug localization based on bug reports. In *Proceedings of the 34th International Conference on Software Engineering*. IEEE, 14–24.
- [80] Guido Zuccon, Bevan Koopman, Peter Bruza, and Leif Azzopardi. 2015. Integrating and evaluating neural word embeddings in information retrieval. In *Proceedings of the 20th Australasian Document Computing Symposium*. ACM, 12.