

Keeping Deep Learning Models in Check: A History-Based Approach to Mitigate Overfitting

Hao Li*, Gopi Krishnan Rajbahadur†, Dayi Lin†, Cor-Paul Bezemer*, and Zhen Ming (Jack) Jiang‡

*University of Alberta. {li.hao, bezemer}@ualberta.ca,

†Centre for Software Excellence, Huawei Canada. {gopi.krishnan.rajbahadur1, dayi.lin}@huawei.com,

‡York University. zmjiang@cse.yorku.ca

Abstract—In software engineering, deep learning models are increasingly deployed for critical tasks such as bug detection and code review. However, overfitting remains a challenge that affects the quality, reliability, and trustworthiness of software systems that utilize deep learning models. Overfitting can be (1) prevented (e.g., using dropout or early stopping) or (2) detected in a trained model (e.g., using correlation-based approaches). Both overfitting detection and prevention approaches that are currently used have constraints (e.g., requiring modification of the model structure, and high computing resources). In this paper, we propose a simple, yet powerful approach that can both detect and prevent overfitting based on the training history (i.e., validation losses). Our approach first trains a time series classifier on training histories of overfit models. This classifier is then used to detect if a trained model is overfit. In addition, our trained classifier can be used to prevent overfitting by identifying the optimal point to stop a model’s training. We evaluate our approach on its ability to identify and prevent overfitting in real-world samples. We compare our approach against correlation-based detection approaches and the most commonly used prevention approach (i.e., early stopping). Our approach achieves an F1 score of 0.91 which is at least 5% higher than the current best-performing non-intrusive overfitting detection approach. Furthermore, our approach can stop training to avoid overfitting at least 32% of the times earlier than early stopping and has the same or a better rate of returning the best model.

Index Terms—Software engineering for AI, AI for Software Engineering, Overfitting, Training history, Deep learning

I. INTRODUCTION

The use of Deep Learning (DL) models in software engineering (SE) research and software products has been skyrocketing over the past decade. For instance, DL techniques have been used for automated bug detection [24], code review [61], and software testing [73]. These applications underscore the importance and ubiquity of DL in modern SE.

Overfitting is one of the fundamental issues that plagues DL models [33, 49, 67, 79, 85]. A DL model can be considered overfitting if the model fits just the training data instead of learning the target hypothesis [79]. An overfit model increases the risk of inaccurate predictions, misleading feature importance, and wasted resources [27].

Currently, the problem of overfitting is addressed in SE studies that use DL models by either (1) preventing it from happening in the first place or (2) detecting it in a trained DL model [79, 85]. Overfitting prevention approaches include early stopping [48], data augmentation [64], regularization [35], and modifying the DL model by adding dropout




Fig. 1: Examples of overfit and non-overfit training histories.

layers [69] or batch normalization [31]. However, many of these approaches are intrusive and require modifying the data or the model structure and expertise to execute correctly and even then, they may not work. For instance, adding dropout layers, a popularly used overfitting prevention scheme, when set with a lower threshold or when added to the earlier layers may cause unintentional overfitting [41]. Furthermore, even the non-intrusive prevention approaches such as early stopping incur trade-offs between model accuracy and training time [51]. For example, late stopping when using the early stopping approach may improve model accuracy, but it will also increase training time. Conversely, stopping too early could result in a sub-optimal model performance.

Overfitting detection approaches like k-fold cross-validation, training the DL model with noisy data points and observing if the added noise impacts the DL model’s accuracy [87], checking if the hypothesis of the trained model and the data are independent [81] can generally be resource intensive and time consuming. For instance, Xu et al. [84] report that training a DL model to find two semantically linkable questions in StackOverflow takes about 14 hours. If one were to conduct a 5-fold cross-validation to detect if the constructed DL model is overfitting, they would have to invest 70 hours, which might be prohibitive in practice.

In this paper, we introduce *OverfitGuard*, an approach to both detect and prevent overfitting using training histories. Figure 1 illustrates example training histories (i.e., the training and validation losses curves) of an overfit and a non-overfit DL model. The training and validation losses of the overfit model both decrease at the beginning of the training process. Following that, the validation loss increases while the training loss decreases, resulting in a large gap between the training

and validation losses. Such a trend indicates poor generalization of the trained model to new data. Researchers have previously employed training histories for decision-making in areas such as quantitative data acquisition and model selection [8, 9, 29, 45, 46, 70, 76]. Similarly, our approach trains a time series classifier on a simulated dataset of training histories (i.e., labelled validation loss curves over epochs of training) of DL models that overfit the training data. Our trained time series classifier detects overfitting in a trained DL model by examining the validation loss history (captured as part of the training history). In contrast to existing overfitting detection approaches, our approach does not incur additional resources or costs, as the training history (also known as the learning curve) is a natural byproduct of the training process. Furthermore, our approach (i.e., the trained time series classifier) can be used to prevent overfitting based on the validation losses of recent epochs (e.g., the last 20 epochs).

Although our approach is trained on a simulated dataset, we evaluate it on a real-world dataset, collected from papers published in top AI venues within the last 5 years. We gathered the training histories from these papers that are explicitly labelled as overfitting or non-overfitting by the authors as the ground truth. The main contributions of this paper are as follows:

- Our approach outperforms the state-of-the-art by at least 5% in terms of F-score for overfitting detection, achieving an F-score of 0.91.
- Our approach has the capability to prevent overfitting at least 32% earlier than early stopping while maintaining (and often surpassing) the rate of reaching the optimal epoch (i.e., the epoch that yields the best model).
- We provide a replication package [1] containing our trained classifiers and labelled training histories which can be directly used by other researchers.

Paper organization. This paper is organized as follows. Section II provides background information about our study. Section III gives an overview of related work. Section IV introduces existing approaches for detecting and preventing overfitting. Section V describes the design of our study, while Section VI provides detailed information about our experimental setup. Sections VII and VIII present the results of our study. Section IX discusses potential threats to the validity of our study. Finally, Section X concludes the paper.

II. BACKGROUND

This section provides an introduction to the concepts of the training history in DL and time series classification.

A. Leveraging training history in DL

Training history, also known as the empirical learning curve [46], provides valuable insights into a DL model’s learning progress and performance throughout the training process. The training history stores a record of metrics during the training process, which are usually recorded in each training iteration or epoch (as shown in Figure 1). Training loss and validation loss are commonly used metrics in training histories.

TABLE I: Studied time series classifiers.

Classifier	Description
KNN-DTW*	Uses K-Nearest Neighbors [25] with Dynamic Time Warping [17] as the distance metric to classify time series data
HMM-GMM	Uses Hidden Markov Model for modelling time series data and Gaussian Mixture Model as the emissions probability density [23, 32]
TSF	Uses a random forest [6] for time series data using an ensemble of time series trees [16]
TSBF	Time Series Bag-of-Features [5] extracts features based on the bag-of-features approach [21] to create a random forest
SAX-VSM	Symbolic Aggregate approXimation [38] converts the data into symbolic representations and Vector Space Model [50, 57] transforms them into vectors to calculate similarity for classification
BOSSVS	Bag-of-SFA Symbols in Vector Space [59] is similar to SAX-VSM but use SFA [60] to transform the data instead of SAX

* KNN-DTW handles variable-length time series data.

Typically, a dataset is divided into training, validation, and test sets. While the training loss reflects how well the DL model learns from the training data during the training process, the validation loss is evaluated based on the validation data, which serves as a proxy for evaluating the model’s performance on unseen data. After training is completed, the trained DL model’s performance is evaluated using the test set that has not been exposed to the model.

Researchers and developers can identify potential issues such as overfitting or underfitting by analyzing the training histories. For example, overfitting is often observed as an increasing divergence between training loss and validation loss over time (as illustrated in Figure 1a). In this paper, we propose an approach that leverages training history to automatically detect and prevent overfitting in DL models.

B. Time series classification

Time series data consists of data points recorded over time, with each point being associated with a specific timestamp and its corresponding value. Time series classification is a machine learning task that aims to categorize time series data into predefined classes. In the context of our study, we consider the training history of a DL model as time series data, as the task of identifying whether a DL model is overfitting based on its training history can be framed as a time series classification problem. Since there has been no prior systematic research on time series classifiers specifically designed for training histories, we have selected six classifiers (shown in Table I) that have been reported as baselines or state-of-the-art in prior studies [2, 77, 78, 82]. These classifiers were chosen due to their demonstrated effectiveness in various time series classification tasks and their potential applicability to the overfitting detection problem in DL.

III. RELATED WORK

A. Mitigating overfitting in SE

Overfitting poses a significant risk to the trustworthiness of software systems and the research studies that employ DL models. SE researchers typically use either overfitting detection or prevention methods to mitigate the problem of overfitting. Among the overfitting prevention methods, dropout is the most commonly adopted approach [79, 85]. Researchers have used dropout in various domains such as code generation [39, 40], logging locations recommendation [37], and comment completion [14, 80]. Regularization is another prominent overfitting prevention strategy that has been used in SE [4, 30, 84, 86], which requires adding another layer to the model structure as well. For example, Zampetti et al. [86] employed L2-norm regularization in training CNN and RNN models to manage self-admitted technical debt in source code.

Early stopping is another frequently used technique to prevent overfitting during the training process of DL models [13, 28, 62]. For example, Shi et al. [62] utilized early stopping when training a deep Siamese network to identify hidden feature requests posted in chat messages by developers. Other techniques like data augmentation [3, 19, 43] and data balancing [53, 75] are also employed to address overfitting. For instance, Bao et al. [3] developed a CNN-based image classification model to filter out non-code and noisy-code frames from programming screencasts. To enhance training data diversity, they employed data augmentation techniques such as rotation, scaling, translation, and shearing.

In terms of overfitting detection approaches, Zhang et al. [87] propose the perturbation validation (PV) assessment to determine whether a DL model fits the training data properly (i.e., ensure that it is neither overfitting nor underfit). Alternatively, some detection approaches check the hypothesis that the trained DL model and the data are independent. For example, Werpachowski et al. [81] check the hypothesis by comparing the test error with the estimated test error based on adversarial examples of the test set.

Another popular approach towards detecting overfitting is to use model validation approaches. Tantithamthavorn et al. [72] evaluated 12 model validation techniques specifically for defect prediction models and concluded that out-of-sample bootstrap emerges as the least biased and most stable technique that helps detect overfit. Damiani and Ardagna [15] introduced a framework that validates DL models against desired non-functional properties and statistically monitors the model output. Straub [71] extended this by utilizing randomly generated expert networks for model validation that focuses on performance characteristics.

Other than these approaches, Smith et al. [67] discussed the role of human factors in overfitting and provided a comparative analysis between automated patches and human-written fixes. They reported that overfitting is not solely a machine-induced problem and suggested focusing on the contributing factors like test suite coverage and requirements-based testing. Xin and Reiss [83] proposed a classification technique that dif-

ferentiates overfitting patches based on semantic differences. Nilizadeh et al. [49] utilized formal verification to evaluate the degree of overfitting and identified the challenges posed by program complexity and numeric issues.

To the best of our knowledge, both the overfitting prevention and detection methods used in both SE studies and in practice fall prey to several key concerns. The overfitting prevention methods, typically require significant expertise to execute correctly and are intrusive (for instance, they may require one to modify the data or the model structure). Overfitting detection approaches typically require retraining of the DL model multiple times, which may be very costly in practice [20] and simple methods like early stop may stop the training of DL model sub optimally. Our work addresses these gaps by introducing a history-based approach that serves the dual purpose of detecting and preventing overfitting in a non-intrusive manner without any need for DL model retraining.


B. Leveraging training history to improve DL software quality

While the machine learning community has leveraged training histories (i.e., learning curves) for different tasks, the SE community has seldom used training history to enhance software quality. Mohr and van Rijn [46] conduct a survey on approaches based on learning curves for decision-making in DL domains, such as data acquisition, early stopping, and model selection. They also propose an approach called Learning Curve Cross-Validation (LCCV) [45] that iteratively increases the number of training examples used for training to select the best model from the candidates. Training histories can also be used to evaluate the trained classifiers: van Rijn et al. [76] propose an approach that recommends classifiers for a given dataset based on training histories using loss time curves. Moreover, Hoiem et al. [29] investigate the use of training histories to evaluate design choices for DL models, such as pretraining, architecture, and data augmentation.


In this paper, we propose *OverfitGuard*, an approach that utilizes time-series classifiers to detect and prevent overfitting by analyzing the training history of DL models. Our approach aims to enhance the software quality of DL systems by improving the quality of the DL models themselves. *OverfitGuard* is one of the first approaches that leverages training history to enhance the software quality of DL systems. A related work by Tokui et al. [74] introduces an approach called *NeuRecover* which also leverages training history to improve the quality of DL models, particularly for safety-critical applications. *NeuRecover* identifies the model parameters that need to be modified (i.e., repaired) by analyzing the training history to address specific failure types. Our approach shares similarities with *NeuRecover*, but focuses on the detection and prevention of overfitting, which is important for maintaining the software quality of DL systems.

IV. EXISTING APPROACH

In this section, we introduce the existing approaches that we use as baseline approaches to compare our proposed approach for overfitting detection and prevention in DL models.



(a) Examples of early stopping with the patience of 20 and 40 epochs (stop at the #71 epoch and #140 epoch).



(b) An demonstration of our overfitting prevention approach with a rolling window.

Fig. 2: Early stopping and our approach for overfitting prevention.

A. Overfitting detection

Correlation-based approaches. One approach for detecting overfitting in DL models is to compute the correlation between the training and validation loss. Kronberger et al. [34] propose computing the non-parametric Spearman’s rank correlation coefficient [68] between training and validation fitness in symbolic regression models to detect overfitting. Similar to our approach, correlation-based approaches also detect overfitting based on the training history. Therefore, we select them as a baseline for comparison. In this study, we calculate the correlation metrics between the training and validation loss to detect overfitting in DL models. The underlying principle is straightforward: when no overfitting occurs, the training and validation losses should be correlated, whereas weak correlation implies overfitting.

We determine the presence of overfitting by comparing the computed correlation-based metric between training and validation losses with a predetermined threshold value (determined in Section VI-D). We use three different correlation metrics: Spearman, Pearson [26], and time-lagged Pearson correlation coefficients. Both Spearman and Pearson correlation coefficients are calculated since we do not know whether the relationship between training and validation loss is linear or not. In addition, we compute the Pearson correlation coefficient between a 5-epoch lagged version of the training loss and the validation loss. This approach is inspired by autocorrelation [10], which measures the correlation between a time series data and a time-lagged version of itself.

B. Overfitting prevention

Early stopping. One widely used approach for preventing overfitting is early stopping, which stops training when there is no improvement in a fixed number of epochs (called patience) and returns the epoch which has the lowest validation loss. We choose the widely used TensorFlow implementation,¹ which

¹https://tensorflow.org/api_docs/python/tf/keras/callbacks/EarlyStopping

is also used by PyTorch Ignite.² As shown in Figure 2a, early stopping with a patience parameter of 20 epochs stops at the #70 epoch and returns the #50 with the lowest validation loss since no improvement occurs between epochs #50 and #70. Furthermore, early stopping with larger patience values (e.g., 40 epochs) stops later (at the #140 epoch) but with a lower loss (at the #100 epoch).

Early stopping based on smoothed validation loss curves.

An alternate version of early stopping inspects the moving average of the smoothed validation loss curves [47, 65] to decide when to stop the training process. After stopping the training process, this approach returns the best epoch which has the lowest validation loss (not the smoothed value).

V. OUR APPROACH

Figure 3 shows an overview of our proposed approach. Our approach uses a time series classifier to detect and prevent overfitting. Table I lists the studied time series classifiers. First, we collect a simulated dataset (more details on how we collect the data in Section VI) that contains training histories (i.e., training and validation loss curves, however, we only use the validation loss curves in our approach) with labels indicating whether overfitting occurs in order to train our time series classifier. Second, we train and evaluate each studied time series classifier on all of the training histories of the simulated dataset. Finally, we use the trained time series classifier to perform both overfitting detection and prevention as follows.

Overfitting detection. To detect overfitting in a trained DL model, we first collect its validation losses over the training epochs. We feed this loss to our trained time series classifier to detect whether there is overfitting. However, we cannot directly feed these validation losses to our classifier, since the length of the validation losses might not be of the same length as that of the data used to train these time series classifiers. All the studied time series classifiers, with the exception of

²https://pytorch.org/ignite/generated/ignite.handlers.early_stopping.EarlyStopping

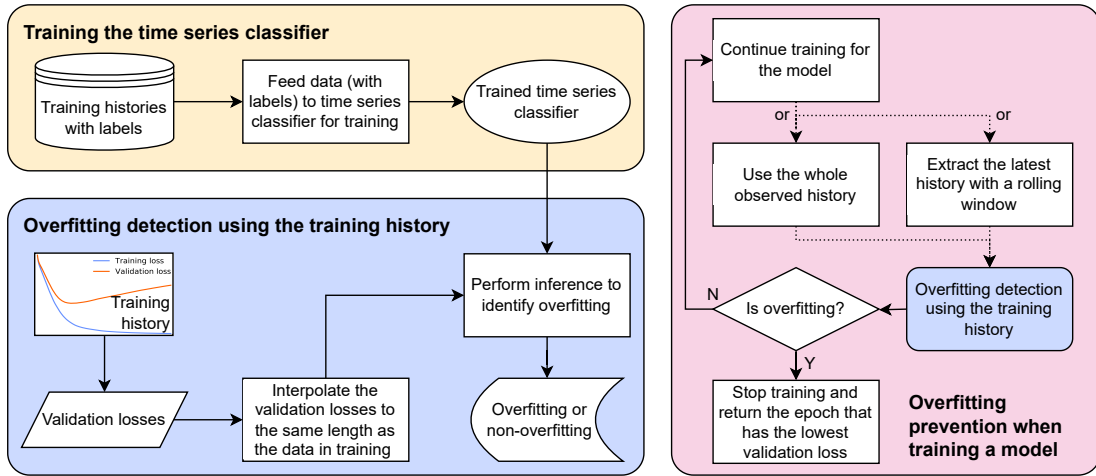


Fig. 3: Our approach for overfitting detection and prevention.

KNN-DTW, expect the length of the inputs used for training and inference to be the same. Therefore, we first linearly interpolate the validation losses of the DL model for which we need to detect overfit to the same length as the training histories used to train the studied time series classifiers. We feed the interpolated validation losses to our trained time series classifier and perform inference to determine if the DL model is overfit.

Overfitting prevention. To prevent overfitting, we feed the training history (i.e., validation loss curve) of a DL model that is being trained to our trained time series classifier during the training process. The history is fed for inference in two different ways: (1) as a rolling window: we extract the latest history in a fixed window size (e.g., the latest 20 epochs), and (2) as the whole observed history (from the first to the latest epochs). Our time series classifier detects if in the fed history overfitting occurs. Similar to overfitting detection, we linearly interpolate the data before feeding it into our model. If there is no overfit occurring, we continue the training and repeat the above procedure until the DL model has finished training. For the rolling window, we move the window by a fixed step size (as shown in Figure 2b) and make another prediction. If our model detects the presence of overfitting in the fed history, we return the lowest validation loss in the observed epochs as the best epoch.

VI. EXPERIMENTAL SETUP

In this section, we introduce the datasets for training and evaluating the studied overfitting detection and prevention approaches, the experiments of our study, and the evaluation metrics for the studied approaches. Figure 4 shows an overview of the experimental setup.

A. Environment setting

We conducted the experiments on an Ubuntu 20.04 operating system with a Linux kernel version of 5.15.0, utilizing

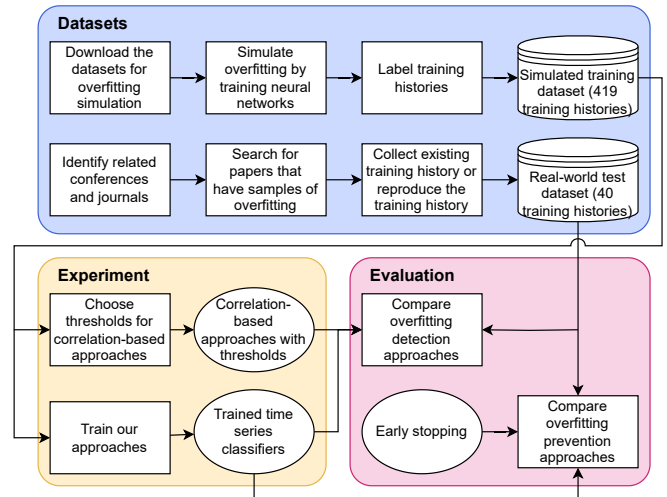


Fig. 4: Overview of the experimental setup.

Python 3.8 and TensorFlow 2.9.0. The hardware configuration for the experiments is detailed below:

- NVIDIA RTX 3090 GPU with 24 GB memory
- CUDA version 10.1.243
- cuDNN version 7.6.5
- Intel(R) Core(TM) i9-11900K CPU with a clock speed of 3.50 GHz
- 64 GB of RAM

B. Simulated training dataset

We create a simulated dataset containing training histories with labels to determine a threshold for correlation-based approaches (see Section IV) and to train our proposed method as described in Section V. We create this simulated dataset by training neural networks of varying model complexities to produce overfitting and non-overfitting samples. The process is as follows:

TABLE II: Information about datasets used to simulate overfitting.

Dataset	Type	#In	#Out	#Examples
building	regression	14	3	4,208
cancer	classification	9	2	699
card	classification	51	2	690
diabetes	classification	8	2	768
flare	regression	24	3	1,066
gene	classification	120	3	3,175
glass	classification	9	6	214
heart	classification	35	2	920
hearta	regression	35	1	920
horse	classification	58	3	364
soybean	classification	82	19	683
thyroid	classification	21	3	7,200

Step 1 – Download the datasets for overfitting simulation.

We download 12 datasets representing real-world problems from the Proben1 [52] benchmark set for training neural networks. These datasets were used by Prechelt [51] to simulate training histories for studying early stopping. We choose these datasets over using SE datasets for two reasons: First, since we use the methodology used by Prechelt [51] to simulate overfitting, we chose to stick with the datasets that they used. Second, irrespective of the domain of the dataset, we assert that how the phenomenon of overfit is represented by training and validation histories will remain the same. Table II provides information about these datasets, which include 3 datasets for regression tasks and 9 datasets for classification tasks. All of these datasets (except the “building” one) were originally collected from the UCI machine learning repository [7], which has been widely used in DL research [22, 36, 56, 63]. Each dataset is pre-partitioned into training, validation, and test sets (50%, 25%, and 25% of the data, respectively) and partitioned three times to generate three distinct permutations, resulting in a total of 36 datasets from Proben1.

Step 2 – Simulate overfitting by training neural networks.

We train neural networks (NNs) with various architectures on the collected 36 datasets. We do so to vary the model complexity which in turn increases the chance of producing an overfitted DL model, following the methodology of Prechelt [51] used in their study. The input/output layer of each NN contains the same number of nodes as the number of input/output coefficients in the respective datasets (see Table II) and rectified linear units (ReLUs) are used for all hidden layers. The structures of the NNs are as follows: (1) 6 one-hidden-layer NNs with 2, 4, 8, 16, 24, or 32 hidden nodes, and (2) 6 two-hidden-layer NNs with hidden nodes (represented as first layer hidden nodes + second layer hidden nodes) of 2+2, 4+2, 4+4, 8+4, 8+8, 16+8. We use the mean square error (MSE) as the loss function for regression problems, and cross entropy as the loss function for classification problems. All problems employ stochastic gradient descent (SGD) as the optimizer. To increase the likelihood of overfitting, we train these 12 neural network architectures on each of the 36 datasets for 1,000 epochs, producing 432 training histories.

Step 3 – Label training histories. To ensure the robustness

of our manual labelling process, we follow the approach outlined by Ding et al. [18]. The first and second authors of this paper independently labelled the 432 data points as either “overfit”, “non-overfit” or “uncertain” and discussed the results. In the first discussion round, the authors reached a 95% agreement (410 data points), with both authors labelling 10 data points as “uncertain” and subsequently eliminating them. In the second round, the authors discussed the remaining 22 disagreements. Following the discussion, we eliminated 3 data points (labelled “uncertain” by both authors) and agreed on the labels for the remaining 19 data points. The final dataset consists of 44 overfit and 375 non-overfit training histories. We share the labelled training histories in our replication package for other researchers to reuse.

C. Real-world test dataset

To evaluate our approach using real-world data, we conducted a survey of papers from conferences and journals to gather examples of overfit and non-overfit DL models.

Step 1 – Identify related conferences and journals. We identify related conferences and journals based on the Computing Research and Education Association of Australasia (CORE³) and China Computer Federation (CCF⁴) ranking systems. Under the CCF A rank, we have 7 conferences and 4 journals in the “Artificial Intelligence” field. Under the CORE A* rank, we have 16 conferences in the “machine learning” and “artificial intelligence” fields and 12 journals in the “artificial intelligence and image processing” field. After merging the results and accounting for overlaps between the two ranking systems, we obtained a final list of 17 conferences and 12 journals.

We collected our real-world data in the artificial intelligence and machine learning domain as opposed to SE domain, because SE studies typically do not report the training histories of the overfit DL models and we required community accepted examples of the training histories of both overfit and non-overfit DL models.

Step 2 – Search for papers that have samples of overfitting. We found 33 full papers (see the replication package [1]) containing the keyword “overfit” (including variations such as “overfitting”) in the title that were published in the selected conferences and journals in the last 5 years. Five of these papers provided samples of overfitting: P2 - Chatterjee and Mishchenko [11]; P4 - Chen et al. [12] P13 - Kim et al. [33]; P17 - Rice et al. [54] and P23 - Singla et al. [66]. Table III lists the papers and the number of collected samples of overfitting (some of them also provide samples of non-overfitting).

Step 3 – Collect existing training history or reproduce the training history. Paper P17 shared the training history, making its replication straightforward. We replicated the other papers that provide overfitting samples to collect the training histories

³<https://www.core.edu.au>

⁴<https://ccf.atom.im/>

TABLE III: Information about collected samples from surveyed papers.

Paper	Labels for the training history in the manuscript	#Overfit	#Non-overfit
P2	"[...] the validation accuracy of nn-random is 9.73% (i.e., close to chance) confirming that it is horribly overfit"	2	0
P4	"We first observe that the robust overfitting prevails in all Baseline cases" "Our methods effectively mitigates the robust overfitting"	3	3
P13	"Figure 4 [...] that is, catastrophic overfitting occurs." "Figure 6 shows that the proposed method also successfully prevents catastrophic overfitting [...]"	N/A*	N/A*
P17	"Figure 24 [...] We see clear robust overfitting for the smaller two options in λ , and find no overfitting but highly regularized models for the larger two options [...]"	20	4
P23	"These results therefore validate our claim that low curvature activations reduce robust overfitting"	4	4

* Cannot reproduce the same results as the paper.

of these samples. We executed the code from the papers with available replication packages (P4, P13, and P23) to generate their training histories. However, we could not replicate the results for paper P13. For paper P2, which did not provide a replication package, we followed the methodology to replicate the results and training history. In total, we collected 29 training histories of overfit DL models and 11 of non-overfit DL models (refer to Table III for details).

D. Experiments

Overfitting detection. We trained the time series classifiers based on the simulated dataset. For each classifier, a grid search with 3-fold cross-validation was performed to tune the hyperparameters based on the simulated dataset. Once the optimal hyperparameters were identified, we proceeded to train each time series classifier using all training histories and labels from the simulated dataset and saved the trained classifier for further use. For correlation-based approaches, we also performed a grid search based on the simulated dataset to select the optimal thresholds (ranging from -1 to 1) that yielded the best F-score.

Overfitting prevention. We reused the trained time series classifiers from the previous step to perform inference during the training process to prevent overfitting. Since the validation loss curve is generally applicable to both classification and regression tasks, we used it for overfitting prevention. We applied our approach to the trained DL models in every 10 epochs (i.e., the step size), with varying rolling window sizes of 20, 40, 60, 80, and 100 epochs. We used early stopping based on the validation loss and set the patience values to range from 5 to 115 epochs. We also applied early stopping based on smoothed validation loss curves generated by a 10-epoch moving average [47, 65].

E. Evaluation

Evaluation metrics for overfitting detection. To evaluate the classification performance of the overfitting detection

approaches, we computed the **precision**, **recall**, and **F-score** for overfitting and non-overfitting samples in the real-world test dataset. In addition, we calculated the **average F-score** to directly compare the classification performance of the studied approaches. To evaluate the time cost associated with training and using the studied approaches, we report the **training time** (in seconds) for each approach on the simulated dataset and the **inference time** (in milliseconds) for the real-world dataset.

Evaluation metrics for overfitting prevention. Ideally, an overfitting prevention approach returns the optimal epoch (i.e., the epoch that yields the best predictive performance for the DL model on the validation set) and stops the training process as early as possible. We define the **optimal rate** of an overfitting prevention approach as the percentage of cases where the optimal epoch is successfully identified. To assess the speed of the approach, we introduce the **delay** metric, which represents the epoch difference between the stopped epoch and the best epoch. For example, a delay of 10 epochs occurs if the prevention approach stops at the 123th epoch while the 113th epoch is the best one. In addition, we report the DL model’s **accuracy** on the validation set when the training process is stopped by the overfitting prevention approach.

VII. RQ1: HOW WELL DOES OVERFITGUARD DETECT OVERFITTING IN TRAINED DL MODELS?

Motivation. Overfitting detection is an important task in DL models since it helps in identifying whether a DL model has learned to perform well on training data but fails to generalize on unseen data. Accurate overfitting detection can assist researchers and developers in making informed decisions regarding model selection, hyperparameter tuning, and other model performance improvements. This research question investigates the performance of our proposed approach for detecting overfitting in trained DL models and compares it with existing correlation-based approaches.

Approach. We use the evaluation metrics introduced in Section VI-E to compare our approach with baseline approaches based on the real-world test dataset. Furthermore, we record the F-score obtained from the 3-fold cross-validation (CV) for our approach based on the simulated training dataset to further analyze the performance of our approach. Since we use the entire simulated training dataset to determine the thresholds (without CV) for correlation-based approaches, we report the F-score for correlation-based approaches based on the whole simulated training dataset.

Results. *Overfit DL models can be detected by inspecting the training history, and our approach using time series classifiers demonstrates better classification performance than the correlation-based approaches for overfitting detection.* Table IV shows that our approach using KNN-DTW, TSBF, and SAX-VSM generalizes well from the simulated dataset to the real-world dataset with the best F-score (0.91), followed by TSF which outperforms the baseline approaches as well. In contrast, HMM-GMM performs poorly on both the simulated training and real-world test datasets. One possible explanation

TABLE IV: Results of the overfitting detection approaches on the simulated dataset (CV F-S: F-score of cross-validation) and real-world dataset (Prec: precision; Rec: recall; F-s: F-score; Avg F-s: average F-score), and the time cost of training the studied approaches on the simulated dataset and performing inference on the real-world dataset (per sample).

Detection approach		Non-overfitting			Overfitting			Avg F-s	CV F-s	Training time (s)	Inference time (ms)
		Prec	Rec	F-s	Prec	Rec	F-s				
Corr. based	Spearman	0.71	0.91	0.80	0.96	0.86	0.91	0.86	0.95	2.461	0.908
	Pearson	0.78	0.64	0.70	0.87	0.93	0.90	0.80	0.92	0.222	0.025
	Autocorr	0.80	0.73	0.76	0.90	0.93	0.92	0.84	0.80	0.233	0.026
Time series (ours)	KNN-DTW	0.79	1.00	0.88	1.00	0.90	0.95	0.91	0.97	0.001	180.512
	HMM-GMM	0.30	0.27	0.28	0.73	0.76	0.75	0.52	0.59	99.751	17.750
	TSF	0.77	0.91	0.83	0.96	0.90	0.93	0.88	0.99	0.311	17.209
	TSBF	0.79	1.00	0.88	1.00	0.90	0.95	0.91	0.99	0.301	31.683
	BOSSVS	0.46	0.91	0.61	0.94	0.59	0.72	0.67	1.00	1.877	19.342
	SAX-VSM	0.83	0.91	0.87	0.96	0.93	0.95	0.91	0.96	0.912	17.474

is that the extracted state models (via HMM) of the training histories do not follow a Gaussian probability distribution. Our approach with BOSSVS correctly identifies all the data in the simulated dataset but performs poorly on the real-world dataset. One reason could be that the extracted bag-of-SFA symbols (BOSS) from the simulated dataset does not generalize to the real-world dataset. In addition, we note that the investigated correlation-based approaches perform reasonably well, with F-scores greater than 0.8. However, our approach outperforms the correlation-based overfitting detection approach by at least 5% on the studied real-world dataset.

The studied time series classifiers are more computationally intensive than correlation-based approaches for inference, yet they are still useful in practice. As shown in Table IV, our approach requires more time for performing inference than the correlation-based approaches. For instance, TSF has the fastest inference time among the classifiers but is around 20 times slower than the Spearman correlation-based approach and around 700 times slower than the other two correlation-based approaches. However, the speed of our approach is not prohibitive in practice since overfitting detection is only executed once after the training is complete. It is also useful to note that the training times of the time series classifiers in our approach are not excessive. For instance, the training times of TSF and TSBF are around 300 milliseconds, and KNN-DTW, our best-performing time series classifier, can finish training in 1 millisecond. However, KNN-DTW requires the longest time for inference which is around 180 milliseconds for a training history. A fast version of DTW [58] with a time complexity of $O(n)$ is used in experiments, but using KNN with DTW is still computationally intensive.

RQ1 Takeaway: Our proposed approach demonstrates better classification performance than correlation-based approaches for detecting overfitting in DL models. Despite the higher computational cost of the time series classifiers used in our approach, their training time and inference time are still practical.

TABLE V: The optimal rate, median delay, and average accuracy of our overfitting prevention approaches using the whole observed history.

Classifier	Optimal rate	Median delay	Average accuracy
KNN-DTW	0.95	43.5	0.42
HMM-GMM	0.18	0.0	0.36
TSF	0.90	35.0	0.42
TSBF	0.83	31.0	0.41
BOSSVS	0.65	21.0	0.41
SAX-VSM	0.33	10.0	0.38

VIII. RQ2: HOW WELL DOES OVERFITGUARD PREVENT OVERFITTING DURING THE TRAINING PROCESS?

Motivation. Another critical part of developing trustworthy and stable DL models is preventing overfitting. An effective overfitting prevention approach allows DL models to generalize better on unseen data while minimizing both training resources and computational costs. This research question evaluates the performance of our proposed approach for preventing overfitting during the training process compared with the frequently used early stopping approach.

Approach. We assess our overfitting prevention approach against the early stopping method (both with and without smoothing loss curves) using the metrics introduced in Section VI-E. To study the difference in delay across overfitting prevention approaches, we performed the Mann-Whitney U test [44] at a significance level of $\alpha = 0.05$ to determine whether the distributions of the delay epochs of early stopping and our approach are significantly different. We also computed Cliff’s delta d [42] effect size to quantify the difference based on the provided thresholds [55].

Results. *Our proposed approach, utilizing KNN-DTW with both rolling window and whole observed history, has a similar or higher optimal rate than early stopping for overfitting prevention.* Other studied classifiers do not perform as well as KNN-DTW for overfitting prevention. As Figure 5 and Table V show, using KNN-DTW with either a rolling window or the whole observed history outperforms early stopping at

TABLE VI: The median delay and average accuracy of early stopping (es) and our overfitting prevention approaches (using a rolling window) with different window sizes (ws).

Prevention approach	Median delay / ws					Average accuracy / ws				
	20	40	60	80	100	20	40	60	80	100
ES	20.0	40.0	60.0	80.0	100.0	0.40	0.42	0.42	0.43	0.43
ES (smoothed)	20.0	43.5	62.0	82.0	97.5	0.39	0.42	0.43	0.43	0.43
KNN-DTW	31.0	27.0	37.5	42.5	45.5	0.42	0.42	0.43	0.43	0.43
HMM-GMM	5.0	6.5	16.5	28.0	41.5	0.37	0.37	0.38	0.39	0.41
TSF	12.5	22.0	31.0	39.5	44.0	0.40	0.40	0.42	0.43	0.43
TSBF	8.5	15.0	25.0	37.0	47.0	0.38	0.39	0.40	0.41	0.42
BOSSVS	7.0	29.0	34.5	48.5	56.5	0.42	0.42	0.43	0.43	0.43
SAX-VSM	4.0	11.0	9.5	16.0	24.0	0.37	0.37	0.38	0.38	0.40




Fig. 5: The optimal rate of our overfitting prevention approach (using a rolling window) and early stopping with different patience values.

identifying the optimal epoch. In particular, our approach with KNN-DTW based on the rolling window has a higher or the same optimal rate as both early stopping approaches when using up to 80 epochs as the patience parameter and window size. For example, our approach with KNN-DTW obtains 78% optimal rate when setting the window size to 20 epochs, while both early stopping approaches achieve less than 50% optimal rate when the patience parameter is set to the same epochs. However, when the patience parameter is greater than 80 epochs, both early stopping approaches can identify almost all of the optimal epochs. The reason is that 90% of the training histories in the real-world dataset have around 200 epochs, hence, a large patience value makes it easy for early stopping to choose the optimal epoch. In addition, Table V shows that our approach with KNN-DTW based on the whole observed history also obtains a higher optimal rate compared to both early stopping approaches. For example, the KNN-DTW classifier achieves a 95% optimal rate with a median delay of 43.5 epochs, whereas early stopping approaches achieve around 85% using the same number (i.e., between 40 to 45 epochs patience in Figure 5).

Our approach using KNN-DTW and a rolling window can stop training a DL model earlier than early stopping (based




Fig. 6: Our approach using KNN-DTW (set the window size as 40 epochs) stops earlier than the early stopping (set the patience parameter as 40 epochs) but both achieve the same optimal epoch.

on both smoothed and non-smoothed validation loss) with the same or higher accuracy. As shown in Table VI, with the same number of epochs for the patience parameter and window size, our approach can save training time (i.e., reducing delay between the stopped epoch and the best epoch) compared to early stopping, except for a window size of 20 epochs. For instance, when setting both the patience parameter and window size to 40 epochs, KNN-DTW and early stopping have the same average accuracy, but KNN-DTW has a median delay of 27 epochs while early stopping has a fixed delay of 40 epochs. The significance test results indicate that the delay difference between KNN-DTW and early stopping is significant (except when using a window size of 20), with a medium to large effect size. Furthermore, early stopping with smoothed loss curves has a median delay of 43.5, which is slower than using the original loss curves with the same patience parameter (40 epochs). Early stopping using the smoothed loss could hurt the performance of early stopping and cannot compete with our approach. In comparison to the delay in early stopping,

the delay between the stopped epoch and the best epoch is at least 32% shorter with our approach using KNN-DTW. Figure 6 provides an example where both early stopping and our approach identify the optimal epoch, but our approach stops 21 epochs earlier than early stopping (which stops with a 40 epochs delay).

Among our two approaches for overfitting prevention, we recommend using KNN-DTW with a rolling window. Although using the whole observed history may achieve a higher optimal rate than using a rolling window for our approach, we note that we can predict the optimal epoch much earlier with the rolling window approach for a very small trade-off in optimal rate (with a similar average accuracy). As shown in Table VI and Figure 5, our approach with KNN-DTW achieves an 83% optimal rate with a median delay of 27 epochs and a 90% optimal rate with a median delay of 37.5 epochs using the window size as 40 and 60 epochs respectively. However, the median delay of KNN-DTW when using the whole observed history is 43.5, while the rolling window approach with a window size of 80 or more epochs can achieve a higher optimal rate (98% vs. 95% accuracy) with a shorter delay (42.5 vs. 43.5 epochs). In summary, we suggest using the rolling window approach since it stops earlier with a relatively small optimal rate drop using a small window (e.g., 40 epochs) and outperforms the whole observed history approach when using a large window size (e.g., 80 epochs).

RQ2 Takeaway: Our proposed approach using KNN-DTW with a rolling window or whole observed history outperforms early stopping for overfitting prevention and can stop training DL models earlier with the same or higher accuracy. Among our two approaches, we recommend using KNN-DTW with a rolling window for early stopping, which achieves a high optimal rate with a shorter delay.

IX. THREATS TO VALIDITY

A. Construct Validity

The construct validity of our approach may be affected by the manual labelling process for the simulated training dataset used in overfitting detection. The definition of overfitting is an abstract concept and may result in ambiguity or disagreements among authors. To mitigate this threat, two authors labelled the training histories independently, achieving a 95% agreement rate. Following this, the authors engaged in multiple rounds of discussions to resolve any disagreements (as detailed in Section VI-B). Despite these efforts, some subjectivity in the labelling process might impact the validity of our results.

Another potential threat to construct validity is the choice of the monitoring metric used in overfitting prevention. Although validation loss is a widely used metric for monitoring the DL model performance during the training process, different DL tasks may require alternative metrics. We conducted additional experiments using classification error (i.e., zero-one loss) for

overfitting prevention, and our approach using KNN-DTW still outperformed early stopping.

B. Internal Validity

Our proposed approach relies on the assumption that overfitting can be detected and prevented through the analysis of DL model training histories. However, certain cases of overfitting may not be captured by examining the training histories alone. For instance, data leakage caused by data augmentation or preprocessing in the entire dataset before data splitting (into training, validation, and test sets) could lead to overfitting, but detecting or preventing it solely by inspecting the training history would be challenging.

C. External Validity

We evaluated our proposed approach using a real-world dataset that contains training histories from top AI venues. However, it is still possible that the approach may not generalize well to all types of DL models or datasets. Secondly, our real-world evaluation is based on only 40 data points (of which 29 training histories belong to overfit models), which might not be enough data points to claim generalizability of our proposed approach. Please note that collecting authoritative examples of overfit training history is very hard since researchers and practitioners typically do not report the training history of models that were overfit. In addition, collecting these data points requires one to replicate the studies that report overfit models, which is a very time and resource intensive task. Hence we were limited to 40 data points for the real-world dataset in our study. However, we invite future research to verify the validity of OverfitGuard using our replication package on their own DL model training histories. In addition, the computational resources required to use the proposed approach for inference could limit its applicability in specific situations. For instance, the increased computational cost may be prohibitive in environments with constrained computational resources, while our approach demonstrates improved performance in overfitting detection and prevention than existing approaches.

X. CONCLUSION AND FUTURE WORK

In this paper, we propose a non-intrusive overfitting detection and prevention approach using time series classifiers trained on the training history of DL models. Our approach (when using the KNN-DTW time series classifier) has (1) better classification performance than correlation-based approaches for overfitting detection, and (2) greater accuracy than early stopping for overfitting prevention with a shorter delay. We evaluate our approach on a real-world dataset of labelled training histories collected from the papers published at top AI venues in the last 5 years. Our approach can be a useful tool for researchers and developers of DL software. We have shared the trained time series classifiers in the replication package for reuse, along with all of the training histories and labels. One limitation of our approach is that our best-performing time series classifier takes longer to perform the

inference required to detect and prevent overfit than the studied baselines. We encourage future work to optimize time series classifiers to enable overfitting detection and prevention in real-time with smaller delays. We also encourage the future work to investigate if adding more real world examples to OverfitGuard's training data or using online training to constantly update OverfitGuard would improve its detection and prevention performance.

REFERENCES

- [1] "Our replication package," <https://github.com/asgaardlab/OverfitGuard>, 2023.
- [2] B. S. Anami and V. A. Bhandage, "A comparative study of suitability of certain features in classification of bharatanatyam mudra images using artificial neural network," *Neural Processing Letters*, vol. 50, no. 1, pp. 741–769, 2019.
- [3] L. Bao, Z. Xing, X. Xia, D. Lo, M. Wu, and X. Yang, "Psc2code: Denoising code extraction from programming screencasts," *ACM Trans. Softw. Eng. Methodol.*, vol. 29, no. 3, Jun. 2020.
- [4] A. Barbez, F. Khomh, and Y.-G. Guéhéneuc, "Deep learning anti-patterns from code metrics history," in *2019 IEEE International Conference on Software Maintenance and Evolution (ICSME)*, 2019, pp. 114–124.
- [5] M. G. Baydogan, G. Runger, and E. Tuv, "A bag-of-features framework to classify time series," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 35, no. 11, pp. 2796–2802, 2013.
- [6] G. Biau and E. Scornet, "A random forest guided tour," *Test*, vol. 25, no. 2, pp. 197–227, 2016.
- [7] C. Blake, "UCI repository of machine learning databases," <https://archive.ics.uci.edu/ml/index.php>, 1998.
- [8] J. Bornschein, F. Visin, and S. Osindero, "Small data, big decisions: Model selection in the small-data regime," in *Proceedings of the 37th International Conference on Machine Learning (ICML'20)*, 2020.
- [9] P. Brazdil, J. N. van Rijn, C. Soares, and J. Vanschoren, *Metalearning: Applications to Automated Machine Learning and Data Mining*, 2022.
- [10] P. J. Brockwell and R. A. Davis, *Introduction to time series and forecasting*, 2002.
- [11] S. Chatterjee and A. Mishchenko, "Circuit-based intrinsic methods to detect overfitting," in *Proceedings of the 37th International Conference on Machine Learning*, vol. 119, 13–18 Jul 2020, pp. 1459–1468.
- [12] T. Chen, Z. Zhang, S. Liu, S. Chang, and Z. Wang, "Robust overfitting may be mitigated by properly learned smoothening," in *International Conference on Learning Representations*, 2021.
- [13] M. Choetkiertikul, H. K. Dam, T. Tran, T. Pham, A. Ghose, and T. Menzies, "A deep learning model for estimating story points," *IEEE Transactions on Software Engineering*, vol. 45, no. 7, pp. 637–656, 2019.
- [14] A. Ciurumelea, S. Proksch, and H. C. Gall, "Suggesting comment completions for python using neural language models," in *2020 IEEE 27th International Conference on Software Analysis, Evolution and Reengineering (SANER)*, 2020, pp. 456–467.
- [15] E. Damiani and C. A. Ardagna, "Certified machine-learning models," in *SOFSEM 2020: Theory and Practice of Computer Science: 46th International Conference on Current Trends in Theory and Practice of Informatics, SOFSEM 2020, Limassol, Cyprus, January 20–24, 2020, Proceedings*, 2020, p. 3–15.
- [16] H. Deng, G. Runger, E. Tuv, and M. Vladimir, "A time series forest for classification and feature extraction," *Information Sciences*, vol. 239, pp. 142–153, 2013.
- [17] H. Ding, G. Trajcevski, P. Scheuermann, X. Wang, and E. Keogh, "Querying and mining of time series data: Experimental comparison of representations and distance measures," *Proc. VLDB Endow.*, vol. 1, no. 2, pp. 1542–1552, Aug. 2008.
- [18] Z. Ding, J. Chen, and W. Shang, "Towards the use of the readily available tests from the release pipeline as performance tests: Are we there yet?" in *Proceedings of the ACM/IEEE 42nd International Conference on Software Engineering (ICSE '20)*, 2020, pp. 1435–1446.
- [19] S. Fakhoury, V. Arnaudova, C. Noiseux, F. Khomh, and G. Antoniol, "Keep it simple: Is deep learning good for linguistic smell detection?" in *2018 IEEE 25th International Conference on Software Analysis, Evolution and Reengineering (SANER)*, 2018, pp. 602–611.
- [20] W. Fu and T. Menzies, "Easy over hard: A case study on deep learning," in *Proceedings of the 2017 11th joint meeting on foundations of software engineering*, 2017, pp. 49–60.
- [21] Z. Fu, G. Lu, K. M. Ting, and D. Zhang, "Music classification via the bag-of-features approach," *Pattern Recognition Letters*, vol. 32, no. 14, pp. 1768–1777, 2011.
- [22] S. Gadde, A. Lakshmanarao, and S. Satyanarayana, "Sms spam detection using machine learning and deep learning techniques," in *2021 7th International Conference on Advanced Computing and Communication Systems (ICACCS)*, vol. 1, 2021, pp. 358–362.
- [23] J.-L. Gauvain and C.-H. Lee, "Maximum a posteriori estimation for multivariate gaussian mixture observations of markov chains," *IEEE Transactions on Speech and Audio Processing*, vol. 2, no. 2, pp. 291–298, 1994.
- [24] Q. Hanam, F. S. d. M. Brito, and A. Mesbah, "Discovering bug patterns in JavaScript," in *Proceedings of the 2016 24th ACM SIGSOFT International Symposium on Foundations of Software Engineering (FSE 2016)*, 2016, pp. 144–156.
- [25] D. J. Hand, "Principles of data mining," *Drug safety*, vol. 30, no. 7, pp. 621–622, 2007.
- [26] J. Hauke and T. Kossowski, "Comparison of values of Pearson's and Spearman's correlation coefficients on the same sets of data," *Quaestiones Geographicae*, vol. 30, no. 2, pp. 87–93, 2011.
- [27] D. M. Hawkins, "The problem of overfitting," *Journal of chemical information and computer sciences*, vol. 44, no. 1, pp. 1–12, 2004.
- [28] T. Hoang, H. Khanh Dam, Y. Kamei, D. Lo, and N. Ubayashi, "Deepjit: An end-to-end deep learning framework for just-in-time defect prediction," in *2019 IEEE/ACM 16th International Conference on Mining Software Repositories (MSR)*, 2019, pp. 34–45.
- [29] D. Hoiem, T. Gupta, Z. Li, and M. Shlapentokh-Rothman, "Learning curves for analysis of deep networks," in *Proceedings of the 38th International Conference on Machine Learning*, vol. 139, 18–24 Jul 2021, pp. 4287–4296.
- [30] Q. Huang, A. Qiu, M. Zhong, and Y. Wang, "A code-description representation learning model based on attention," in *2020 IEEE 27th International Conference on Software Analysis, Evolution and Reengineering (SANER)*, 2020, pp. 447–455.
- [31] S. Ioffe and C. Szegedy, "Batch normalization: Accelerating deep network training by reducing internal covariate shift," in *International conference on machine learning*, 2015, pp. 448–456.
- [32] S. Ji, B. Krishnapuram, and L. Carin, "Variational bayes for continuous hidden markov models and its application to active learning," *IEEE Transactions on Pattern Analysis and Machine Intelligence*, vol. 28, no. 4, pp. 522–532, 2006.
- [33] H. Kim, W. Lee, and J. Lee, "Understanding catastrophic overfitting in single-step adversarial training," *Proceedings of the AAAI Conference on Artificial Intelligence*, vol. 35, no. 9, pp. 8119–8127, May 2021.
- [34] G. Kronberger, M. Kommenda, and M. Affenzeller, "Overfitting detection and adaptive covariant parsimony pressure for symbolic regression," in *Proceedings of the 13th Annual Conference Companion on Genetic and Evolutionary Computation (GECCO '11)*, 2011, pp. 631–638.
- [35] J. Kukačka, V. Golkov, and D. Cremers, "Regularization for deep learning: A taxonomy," *arXiv preprint arXiv:1710.10686*, 2017.
- [36] V. Kuleshov, N. Fenner, and S. Ermon, "Accurate uncertainties for deep learning using calibrated regression," in *Proceedings of the 35th International Conference on Machine Learning*, vol. 80, 10–15 Jul 2018, pp. 2796–2804.
- [37] Z. Li, T.-H. Chen, and W. Shang, "Where shall we log? studying and suggesting logging locations in code blocks," in *2020 35th IEEE/ACM International Conference on Automated Software Engineering (ASE)*, 2020, pp. 361–372.
- [38] J. Lin, E. Keogh, L. Wei, and S. Lonardi, "Experiencing sax: a novel symbolic representation of time series," *Data Mining and knowledge discovery*, vol. 15, no. 2, pp. 107–144, 2007.
- [39] F. Liu, G. Li, Y. Zhao, and Z. Jin, "Multi-task learning based pre-trained language model for code completion," in *2020 35th IEEE/ACM International Conference on Automated Software Engineering (ASE)*, 2020, pp. 473–485.
- [40] Z. Liu, X. Xia, M. Yan, and S. Li, "Automating just-in-time comment updating," in *2020 35th IEEE/ACM International Conference on Automated Software Engineering (ASE)*, 2020, pp. 585–597.
- [41] Z. Liu, Z. Xu, J. Jin, Z. Shen, and T. Darrell, "Dropout reduces underfitting," *arXiv preprint arXiv:2303.01500*, 2023.
- [42] J. D. Long, D. Feng, and N. Cliff, "Ordinal Analysis of Behavioral

- Data,” in *Handbook of Psychology*, Apr. 2003, ch. 25, pp. 635–661.
- [43] A. Mahadi, K. Tongay, and N. A. Ernst, “Cross-dataset design discussion mining,” in *2020 IEEE 27th International Conference on Software Analysis, Evolution and Reengineering (SANER)*, 2020, pp. 149–160.
- [44] H. B. Mann and D. R. Whitney, “On a test of whether one of two random variables is stochastically larger than the other,” *Annals of Mathematical Statistics*, vol. 18, pp. 50–60, 1947.
- [45] F. Mohr and J. N. van Rijn, “Towards model selection using learning curve cross-validation,” in *8th ICML Workshop on automated machine learning (AutoML)*, 2021.
- [46] F. Mohr and J. N. van Rijn, “Learning curves for decision making in supervised machine learning - A survey,” *CoRR*, vol. abs/2201.12150, 2022.
- [47] K. Molugaram and G. S. Rao, *Statistical Techniques for Transportation Engineering*, Jan. 2017.
- [48] N. Morgan and H. Bourlard, “Generalization and parameter estimation in feedforward nets: Some experiments,” *Advances in neural information processing systems*, vol. 2, 1989.
- [49] A. Nilizadeh, G. T. Leavens, X.-B. D. Le, C. S. Păsăreanu, and D. R. Cok, “Exploring true test overfitting in dynamic automated program repair using formal methods,” in *2021 14th IEEE Conference on Software Testing, Verification and Validation (ICST)*, 2021, pp. 229–240.
- [50] T. Peng, L. Liu, and W. Zuo, “PU text classification enhanced by term frequency-inverse document frequency-improved weighting,” *Concurrency and Computation: Practice and Experience*, vol. 26, no. 3, pp. 728–741, 2014.
- [51] L. Prechelt, “Early Stopping — But When?” in *Neural Networks: Tricks of the Trade: Second Edition*, ser. Lecture Notes in Computer Science, 2012, pp. 53–67.
- [52] L. Prechelt *et al.*, “Proben1: A set of neural network benchmark problems and benchmarking rules,” 1994.
- [53] X. Ren, Z. Xing, X. Xia, D. Lo, X. Wang, and J. Grundy, “Neural network-based detection of self-admitted technical debt: From performance to explainability,” *ACM Trans. Softw. Eng. Methodol.*, vol. 28, no. 3, Jul. 2019.
- [54] L. Rice, E. Wong, and Z. Kolter, “Overfitting in adversarially robust deep learning,” in *Proceedings of the 37th International Conference on Machine Learning*, vol. 119, 13–18 Jul 2020, pp. 8093–8104.
- [55] J. Romano, J. D. Kromrey, J. Coraggio, J. Skowronek, and L. Devine, “Exploring methods for evaluating group differences on the NSSE and other surveys: Are the t-test and Cohen’sd indices the most appropriate choices,” in *annual meeting of the Southern Association for Institutional Research*. Citeseer, 2006, pp. 1–51.
- [56] S. Sajeev, A. Maeder, S. Champion, A. Beilegoli, C. Ton, X. Kong, and M. Shu, “Deep learning to improve heart disease risk prediction,” in *Machine Learning and Medical Engineering for Cardiovascular Health and Intravascular Imaging and Computer Assisted Stenting*, 2019, pp. 96–103.
- [57] G. Salton, A. Wong, and C.-S. Yang, “A vector space model for automatic indexing,” *Communications of the ACM*, vol. 18, no. 11, pp. 613–620, 1975.
- [58] S. Salvador and P. Chan, “Toward accurate dynamic time warping in linear time and space,” *Intell. Data Anal.*, vol. 11, no. 5, pp. 561–580, Oct. 2007.
- [59] P. Schäfer, “Scalable time series classification,” *Data Mining and Knowledge Discovery*, vol. 30, no. 5, pp. 1273–1298, 2016.
- [60] P. Schäfer and M. Höggqvist, “SFA: A symbolic fourier approximation and index for similarity search in high dimensional datasets,” in *Proceedings of the 15th International Conference on Extending Database Technology (EDBT ’12)*, 2012, pp. 516–527.
- [61] O. B. Sghaier and H. Sahraoui, “A multi-step learning approach to assist code review,” in *2023 IEEE International Conference on Software Analysis, Evolution and Reengineering (SANER)*, 2023, pp. 450–460.
- [62] L. Shi, M. Xing, M. Li, Y. Wang, S. Li, and Q. Wang, “Detection of hidden feature requests from massive chat messages via deep siamese network,” in *2020 IEEE/ACM 42nd International Conference on Software Engineering (ICSE)*, 2020, pp. 641–653.
- [63] Q. Shi, R. Katuwal, P. Suganthan, and M. Tanveer, “Random vector functional link neural network based ensemble deep learning,” *Pattern Recognition*, vol. 117, p. 107978, 2021.
- [64] C. Shorten and T. M. Khoshgoftaar, “A survey on image data augmentation for deep learning,” *Journal of big data*, vol. 6, no. 1, pp. 1–48, 2019.
- [65] R. H. Shumway and D. S. Stoffer, *Time Series Analysis and Its Applications: With R Examples*, ser. Springer Texts in Statistics, 2017.
- [66] V. Singla, S. Singla, S. Feizi, and D. Jacobs, “Low curvature activations reduce overfitting in adversarial training,” in *Proceedings of the IEEE/CVF International Conference on Computer Vision (ICCV)*, Oct. 2021, pp. 16423–16433.
- [67] E. K. Smith, E. T. Barr, C. Le Goues, and Y. Brun, “Is the cure worse than the disease? overfitting in automated program repair,” in *Proceedings of the 2015 10th Joint Meeting on Foundations of Software Engineering (ESEC/FSE 2015)*, 2015, pp. 532–543.
- [68] C. Spearman, “The proof and measurement of association between two things,” *The American journal of psychology*, vol. 100, no. 3/4, pp. 441–471, 1987.
- [69] N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever, and R. Salakhutdinov, “Dropout: a simple way to prevent neural networks from overfitting,” *The journal of machine learning research*, vol. 15, no. 1, pp. 1929–1958, 2014.
- [70] B. Strang, P. v. d. Putten, J. N. v. Rijn, and F. Hutter, “Don’t rule out simple models prematurely: A large scale benchmark comparing linear and non-linear classifiers in openml,” in *Advances in Intelligent Data Analysis XVII*, 2018, pp. 303–315.
- [71] J. Straub, “Machine learning performance validation and training using a ‘perfect’ expert system,” *MethodsX*, vol. 8, p. 101477, 2021.
- [72] C. Tantithamthavorn, S. McIntosh, A. E. Hassan, and K. Matsumoto, “An empirical comparison of model validation techniques for defect prediction models,” *IEEE Transactions on Software Engineering*, vol. 43, no. 1, pp. 1–18, 2017.
- [73] Y. Tian, K. Pei, S. Jana, and B. Ray, “Deepest: Automated testing of deep-neural-network-driven autonomous cars,” in *Proceedings of the 40th International Conference on Software Engineering (ICSE ’18)*, 2018, pp. 303–314.
- [74] S. Tokui, S. Tokumoto, A. Yoshii, F. Ishikawa, T. Nakagawa, K. Munakata, and S. Kikuchi, “Neurecover: Regression-controlled repair of deep neural networks with training history,” in *2022 IEEE International Conference on Software Analysis, Evolution and Reengineering (SANER)*, 2022, pp. 1111–1121.
- [75] M. Tufano, C. Watson, G. Bavota, M. D. Penta, M. White, and D. Poshyvanyk, “An empirical study on learning bug-fixing patches in the wild via neural machine translation,” *ACM Trans. Softw. Eng. Methodol.*, vol. 28, no. 4, Sep. 2019.
- [76] J. N. van Rijn, S. M. Abdulrahman, P. Brazdil, and J. Vanschoren, “Fast algorithm selection using learning curves,” in *Advances in Intelligent Data Analysis XIV*, 2015, pp. 298–309.
- [77] O. Varol, E. Ferrara, F. Menczer, and A. Flammini, “Early detection of promoted campaigns on social media,” *EPJ data science*, vol. 6, pp. 1–19, 2017.
- [78] Z. Wang, L. Wang, C. Huang, Z. Zhang, and X. Luo, “Soil-moisture-sensor-based automated soil water content cycle classification with a hybrid symbolic aggregate approximation algorithm,” *IEEE Internet of Things Journal*, vol. 8, no. 18, pp. 14003–14012, 2021.
- [79] C. Watson, N. Cooper, D. N. Palacio, K. Moran, and D. Poshyvanyk, “A systematic literature review on the use of deep learning in software engineering research,” *ACM Trans. Softw. Eng. Methodol.*, vol. 31, no. 2, Mar. 2022.
- [80] B. Wei, Y. Li, G. Li, X. Xia, and Z. Jin, “Retrieve and refine: Exemplar-based neural comment generation,” in *Proceedings of the 35th IEEE/ACM International Conference on Automated Software Engineering (ASE ’20)*, 2021, pp. 349–360.
- [81] R. Werpachowski, A. György, and C. Szepesvári, “Detecting overfitting via adversarial examples,” *Advances in Neural Information Processing Systems*, vol. 32, 2019.
- [82] X. Xi, E. Keogh, C. Shelton, L. Wei, and C. A. Ratanamahatana, “Fast time series classification using numerosity reduction,” in *Proceedings of the 23rd international conference on Machine learning*, 2006, pp. 1033–1040.
- [83] Q. Xin and S. P. Reiss, “Identifying test-suite-overfitted patches through test case generation,” in *Proceedings of the 26th ACM SIGSOFT International Symposium on Software Testing and Analysis (ISSTA 2017)*, 2017, pp. 226–236.
- [84] B. Xu, D. Ye, Z. Xing, X. Xia, G. Chen, and S. Li, “Predicting semantically linkable knowledge in developer online forums via convolutional neural network,” in *2016 31st IEEE/ACM International Conference on Automated Software Engineering (ASE)*, 2016, pp. 51–62.
- [85] Y. Yang, X. Xia, D. Lo, and J. Grundy, “A survey on deep learning for software engineering,” *ACM Comput. Surv.*, vol. 54, no. 10s, Sep. 2022.

- [86] F. Zampetti, A. Serebrenik, and M. Di Penta, "Automatically learning patterns for self-admitted technical debt removal," in *2020 IEEE 27th International Conference on Software Analysis, Evolution and Reengineering (SANER)*, 2020, pp. 355–366.
- [87] J. M. Zhang, E. T. Barr, B. Guedj, M. Harman, and J. Shawe-Taylor, "Perturbed model validation: A new framework to validate model relevance," *CoRR*, vol. abs/1905.10201, 2019.