





SPERT: Reinforcement Learning-Enhanced Transformer Model for Agile Story Point Estimation

Waleed Younas ^{*}, Rui Chen [†], Jing Zhao [‡] and Tahreem Iqbal [§]

*School of Software Technology
Dalian University of Technology
Dalian, China*

**waleed.unas@gmail.com*


†chen-ray-cr@hotmail.com

‡zhaoj9988@dlut.edu.cn

§tahreem.waleed@gmail.com

Mohamed Sharaf 

*Industrial Engineering Department, College of Engineering
King Saud University, Riyadh, Saudi Arabia
mfsharaf@ksu.edu.sa*

Azhar Imran 

*Department of Creative Technologies
Air University Islamabad
Islamabad, Pakistan
azhar.imran@au.edu.pk*

Received 23 October 2024

Revised 18 November 2024

Accepted 21 November 2024

Published 6 February 2025

Story point estimation is a key practice in Agile project management that assigns effort values to user stories, helping teams manage workloads effectively. Inaccurate story point estimation can lead to project delays, resource misallocation and budget overruns. This study introduces Story Point Estimation using Reinforced Transformers (SPERT), a novel model that integrates transformer-based embeddings with reinforcement learning (RL) to improve the accuracy of story point estimation. SPERT utilizes Bidirectional Encoder Representations from Transformers (BERT) embeddings, which capture the deep semantic relationships within user stories, while the RL component refines predictions dynamically based on project feedback. We evaluate SPERT across multiple Agile projects and benchmark its performance against state-of-the-art models, including SBERT-XG, LHC-SE, Deep-SE and TF-IDF-SE. Results demonstrate that SPERT outperforms these models in terms of Mean Absolute Error (MAE), Median Absolute Error (MdAE) and Standardized Accuracy (SA). Statistical analysis using Wilcoxon tests and

*Corresponding author.

A12 effect size confirms the significance of SPERT's performance, highlighting its ability to generalize across diverse projects and improve estimation accuracy in Agile environments.

Keywords: Agile project management; story point estimation; transformers; BERT; RL; machine learning.

1. Introduction

Agile project management relies on accurate story point estimation to allocate resources effectively, yet traditional models struggle to keep pace with the evolving and dynamic nature of Agile environments [1, 2]. Since the introduction of the Agile methodology [3], effort estimation has become a critical component of software project management. Accurate estimates are necessary for planning, budgeting and resource allocation. Without precise estimations, projects risk budget overruns, delays or even cancellations [2]. For example, a report by the Standish Group reveals that over 50% of software projects suffer from time and cost overruns, with 31% being canceled before completion [2]. A joint study by McKinsey and the University of Oxford further underscores this, showing that large software projects run 66% over budget and are delivered 33% later than planned [4]. These statistics highlight the importance of effective effort estimation for project management, mitigating risks and increasing project success.

Unlike existing static models, Story Point Estimation using Reinforced Transformers (SPERT) leverages Bidirectional Encoder Representations from Transformers (BERT) embeddings alongside an reinforcement learning (RL) mechanism that refines predictions in real-time. This dual approach not only captures the contextual depth of user stories but also adapts to project feedback, addressing limitations found in traditional machine learning (ML)-based estimation models. Traditionally, effort estimation methods were used in linear methodologies like the Waterfall model, where estimations were performed at the start of the project and applied to large-scale modules. Methods such as The Constructive Cost Model (COCOMO) [5] and Function Point Analysis [6] were widely adopted, relying on predefined metrics and historical project data. These models worked reasonably well in structured environments but struggled when applied to modern, flexible methodologies like Agile, where requirements change frequently, and effort estimation must be revisited regularly [7]. This contrast demonstrates the need for adaptive estimation approaches that can accommodate the iterative nature of Agile projects.

Agile methodologies have fundamentally transformed the software development landscape by promoting iterative development, flexibility and frequent delivery through smaller tasks known as user stories [1]. In Agile, the focus shifts from long-term project-wide estimates to more granular, sprint-based estimates, introducing unique challenges in effort estimation. Agile projects must accommodate volatile requirements, customer feedback and evolving business goals, all of which make accurate effort estimation challenging [8].

One of the major challenges in Agile project management is dealing with volatility. Agile projects frequently experience changing requirements and unforeseen customer feedback during development, which directly impact the estimation of story points [8–10]. A typical example is when a feature is reprioritized midway through a sprint due to a sudden shift in business strategy. This requires developers to revisit their initial effort estimates. This volatility in requirements demands a flexible approach to story point estimation that can dynamically adapt to these changes. Traditional models like COCOMO, which rely on static project data, are not well-suited for such dynamic environments [5].

To address these challenges, several models have been proposed, including Sentence-BERT with XGBoost (SBERT-XG) [11], Lightweight Hybrid Classifier for Software Effort Estimation (LHC-SE) and Deep Learning-based Software Effort Estimation (Deep-SE). SBERT-XG combines Sentence-BERT text embeddings with the XGBoost algorithm to predict story points, capturing the semantic context within user stories more effectively [12, 13]. However, while it improves upon basic text representation models, SBERT-XG is limited in its ability to handle long-term dependencies in textual data. Similarly, LHC-SE employs a hybrid ML approach to effort estimation, leveraging clustering techniques to group similar user stories and improve prediction accuracy [14]. Although this approach helps in reducing estimation variance, it still lacks the flexibility required in highly dynamic Agile environments. Deep-SE utilizes deep learning techniques, such as neural networks, to model complex relationships between story points and user story features [15]. However, despite their promise, deep learning models often require large datasets for training and can be prone to overfitting in smaller Agile projects, limiting their generalization across diverse scenarios [16].

Furthermore, traditional Natural Language Processing (NLP) techniques such as Term Frequency-Inverse Document Frequency (TF-IDF) [17] have been employed in earlier attempts to automate story point estimation. While TF-IDF can provide a basic representation of user stories by encoding word frequency, it struggles to capture deeper semantic relationships and contextual dependencies that are essential for understanding the complexity of Agile tasks. In contrast, more advanced NLP models like BERT have revolutionized the field of natural language understanding by capturing both syntactic and semantic nuances in text [18]. BERT's ability to process bidirectional context (i.e. considering the words before and after a token) makes it particularly suited for analyzing user stories, which often contain complex and context-dependent language. For instance, the phrase “optimize database performance” may imply different levels of complexity depending on the surrounding context, which BERT can capture effectively [19, 20].

While these models offer significant improvements, their primary limitation lies in their static nature — none of them are designed to continuously learn and adapt as project conditions evolve. This is where RL comes into play. RL offers a powerful solution to this problem by allowing models to adjust their behavior based on real-time feedback. RL has been successfully applied in various domains, such as game

playing, robotics and financial forecasting [21]. In the context of Agile software development, RL can help story point estimation models dynamically refine their predictions as project requirements, team dynamics or customer feedback change during development. Proximal Policy Optimization (PPO), in particular, is an efficient RL algorithm known for its stability and robustness in learning policies for continuous tasks [22, 23]. PPO enables the model to adjust story point predictions in response to the feedback it receives, ensuring that it stays adaptable throughout the project's lifecycle [24].

In this paper, we introduce SPERT, a novel model that combines the strengths of BERT-based embeddings with the adaptability of RL. SPERT is specifically designed to address the limitations of existing models in handling dynamic Agile environments. By leveraging pre-trained transformer embeddings, SPERT captures the rich semantic information in user stories, while its RL component ensures that the model continuously improves its predictions based on project feedback. Unlike traditional models like COCOMO, which rely on static historical data, SPERT offers a flexible, data-driven approach that adapts to the evolving conditions of Agile projects.

Our contributions are as follows:

- (1) **Automated Effort Estimation:** We propose a model that automates effort estimation using transformer-based NLP techniques combined with RL, offering a scalable solution for Agile projects.
- (2) **Improved Prediction Accuracy:** By leveraging BERT embeddings and RL, SPERT achieves higher accuracy in story point prediction compared to models like SBERT-XG, LHC-SE and Deep-SE.
- (3) **Integration of NLP Techniques:** SPERT employs advanced NLP methods, such as BERT, to capture the semantic structure of user stories, leading to more reliable effort estimation.
- (4) **Comprehensive Evaluation:** Through extensive experimentation across multiple Agile projects, we demonstrate the effectiveness of SPERT in reducing prediction error and improving generalization across different projects.

The remainder of this paper is organized as follows. Section 2 reviews related work, including traditional and ML-based effort estimation methods. Section 3 describes the dataset used in our experiments. Section 4 outlines the proposed methodology, detailing the architecture of SPERT and its integration of BERT embeddings and RL. Section 5 explains the experimental setup and evaluation metrics. Section 6 elaborates on the model training process. Section 7 presents the evaluation methodology, while Sec. 8 discusses the experimental results. Finally, Sec. 9 concludes the paper with suggestions for future research.

2. Related Work

Software effort estimation (SEE) has evolved over the decades from traditional models to state-of-the-art ML and NLP methods. While older methods like

COCOMO and Function Point Analysis [5, 6] laid the groundwork for early effort estimation, they are ill-suited for modern Agile environments, where project scope changes frequently. These traditional models relied on predefined metrics and upfront estimation, making them less adaptable to Agile's dynamic nature [7].

In recent years, the focus has shifted toward developing ML-based approaches that automate the process of story point estimation. These methods aim to improve estimation accuracy by learning patterns from historical data, thereby reducing reliance on manual input and expert judgment [25]. This section divides the related work into three key areas: (A) ML-based effort estimation, (B) NLP techniques in story point estimation and (C) RL in effort estimation. A fourth subsection (D) is the comparison of existing baseline models and discusses current models used for comparison in this study.

2.1. *ML-based effort estimation*

Early ML models, such as decision trees, support vector machines (SVMs) and linear regression, have shown potential in automating effort estimation [24, 26, 27]. These models aim to predict story points based on historical data, improving consistency over expert-driven estimates. However, the manual feature engineering required by these models often results in biased or incomplete feature sets, limiting their scalability and adaptability to new projects. Moreover, these traditional ML models tend to overlook the sequential nature of Agile stories, which contain rich temporal information [15, 18].

In response to these limitations, deep learning-based methods have emerged as a more powerful tool for effort estimation. Deep neural networks (DNNs) can automatically extract features from raw data, reducing the need for manual engineering [16, 29]. Despite their advantages, deep learning models often require large datasets and are prone to overfitting in small data environments. Furthermore, the “black box” nature of these models can make them difficult to interpret, posing challenges in explaining predictions to stakeholders [30].

2.2. *NLP techniques in story point estimation*

Recent advancements in NLP techniques have made significant contributions to story point estimation by enhancing the ability to process and understand textual data. Traditional NLP methods, such as Bag-of-Words (BoW) and TF-IDF, represent text in a way that overlooks the dependencies between words, resulting in a loss of context and meaning [31]. Although these methods have been used in early ML-based story point estimation models, their inability to capture semantic relationships limits their effectiveness in Agile environments, where user stories are often complex and context-dependent [32].

More recently, deep learning-based NLP methods have taken the forefront. Long Short-Term Memory (LSTM) networks, which capture the sequential dependencies

in text, have shown promise in improving story point estimation accuracy [16, 33]. However, LSTMs struggle with modeling long-range dependencies, making them less suitable for complex user stories with deep contextual relationships [34].

The introduction of transformer-based models like BERT has revolutionized NLP by capturing both long-range dependencies and the nuanced context of text. BERT leverages pre-trained embeddings, which provide richer, more robust representations of user stories compared to earlier methods [18, 19]. In the context of Agile projects, BERT's ability to model the relationships between different parts of a user story is particularly beneficial, as it allows for more accurate story point predictions [35].

2.3. *RL in effort estimation*

RL is an emerging approach in effort estimation, enabling models to continuously learn and adapt to changing project environments. Unlike traditional models, which rely on static data, RL-based approaches learn policies that dynamically adjust predictions based on feedback from previous outcomes [36, 37]. This adaptability makes RL particularly suitable for Agile environments, where project conditions frequently evolve, and story complexities can shift over time [38, 39].

PPO is one of the leading RL algorithms, has been shown to be effective in learning policies for continuous tasks such as effort estimation. PPO's stability and efficiency make it ideal for integrating RL into story point estimation models [22, 40]. While research on combining RL and NLP for story point estimation is still limited, studies have begun exploring how these techniques can work together to provide more accurate and adaptive estimates [41, 42].

2.4. *Comparison of existing baseline models*

Several state-of-the-art models have been developed for Agile story point estimation, each leveraging different techniques to tackle the challenges posed by dynamic project environments. SBERT-XG, for example, integrates Sentence-BERT embeddings with the XGBoost algorithm for regression tasks [43]. While effective in capturing semantic context, its reliance on XGBoost limits its adaptability to evolving Agile stories. LHC-SE uses a hybrid clustering approach to group similar stories for more accurate estimation [39]. However, its static clustering approach struggles with projects, where stories frequently change. Deep-SE applies deep learning techniques to capture complex relationships between user stories and their associated effort, but is prone to overfitting in small datasets [15, 44].

In contrast, SPERT, the model we propose, combines the contextual understanding of transformer-based models like BERT with the adaptive learning capabilities of RL. This hybrid approach enables SPERT to offer more accurate and flexible predictions in Agile environments, outperforming the aforementioned models in both generalization and adaptability [18, 22].

3. Dataset

In this study, we utilized the publicly available dataset provided by Choetkiertikul *et al.* [16], which consists of 23,313 user stories from 16 different open-source projects across nine repositories. This dataset serves as a robust foundation for training and evaluating story point estimation models in Agile environments. The dataset includes a variety of project types, such as mobile development, cloud platforms and educational systems, making it suitable for testing model generalizability across domains.

The repositories in the dataset include well-known projects such as Apache Mesos, Appcelerator Studio, Aptana Studio and JIRA Software, among others. Each issue or user story contains key attributes, including project ID, title, description and the assigned story point estimate. These attributes provide sufficient context for analyzing the relationship between user stories and their associated effort, which is crucial for accurate story point estimation.

3.1. Story point distribution

The story points in the dataset vary across projects, with a wide range of values representing the complexity and required effort for each task. As shown in Fig. 1, the majority of user stories are assigned relatively low story points, with story points often assigned values of 1 or 3. These smaller story point values align with Agile practices, where tasks are broken down into manageable units to facilitate faster iteration and completion.

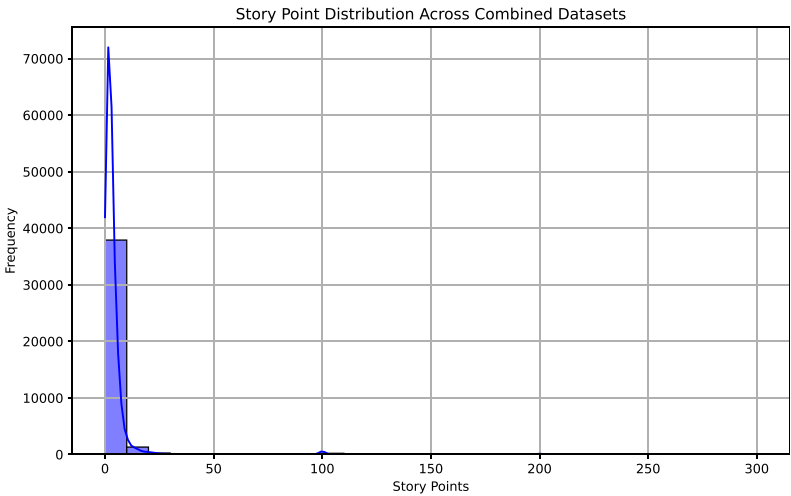


Fig. 1. Story point distribution across combined datasets.

Although specific sprint velocity data are unavailable in this dataset, velocity can be inferred by summing the story points completed over multiple sprints. The story

point distribution aids in estimating team performance, while progress analysis over time can guide future project planning and resource allocation. As Agile methodologies emphasize iterative development and regular feedback, understanding velocity trends is crucial for refining story point estimates and ensuring timely delivery.

Through the insights gained from this dataset, we can simulate how Agile teams adjust to evolving project conditions and improve their estimation accuracy over time.

3.2. *Cross-project dataset diversity*

The dataset includes a diverse set of projects from multiple domains, such as mobile development and educational systems. This diversity makes it ideal for cross-project evaluation, allowing us to assess the model's ability to generalize across different types of Agile projects. The variation in user story complexity and differences in project domains helps evaluate the robustness of story point estimation models like SPERT.

The dataset covers a broad spectrum of Agile software development scenarios, with projects including both feature development and bug fixes. This diversity in tasks and story point distribution enables the development and testing of models that can handle different project dynamics and user story complexities.

3.3. *Example of user stories*

To provide context for how story points are assigned, consider the following example from the Apache Mesos project, illustrated in Fig. 2:

Title: “Changing Theme in Aptana Studio 3 makes selection color/colour opaque — code/text cannot be read/seen”

Description: Changing the theme in Aptana Studio 3 via the round color/colour wheel with a drop-down button makes selection color/colour opaque, i.e. loses transparency. Selected code/text cannot be read/seen. Quitting out of Aptana Studio 3 and re-loading fixes the problem. See attachment for an example.


This issue was assigned a story point estimate of 5, indicating its moderate complexity. The detailed description allows the team to estimate the effort involved in fixing the problem, providing a clear example of how Agile teams use story points to allocate resources and prioritize tasks.

3.4. *Revised dataset statistics*

The key descriptive statistics for the dataset are summarized in Table 1, including the number of issues, minimum and maximum story points, mean, median and standard deviation of the story points. The diversity in story point values across the

Projects /  Apache /  Add epic /  SCRUM-3

Changing Theme in Aptana Studio 3 makes selection color / colour opaque - code/text can not be read / seen

 Attach  Add a child issue  Link issue  

Priority  Medium

Description

Changing Theme in Aptana Studio 3 via round color/colour wheel with drop down button makes selection color / colour opaque, i.e. loses transparency. Selected code/text can not be read / seen. Quitting out of Aptana Studio 3 and then re-loading fixed problem. See attachment from head of a python script after changing theme

Assignee  martha.rollin
[Assign to me](#)

Story point estimate  5

Sprint [SCRUM Sprint 46](#)

Fig. 2. Example of a user story with story points in the Apache Mesos project.

Table 1. Revised descriptive statistics of the story point dataset.

Repo	Project	Abb.	# Issues	Min SP	Max SP	Mean SP	Median SP	Std SP
Apache	Apache Mesos	ME	1680	1	40	3.09	3	2.42
	Usergrid	UG	482	1	8	2.85	3	1.40
Appcelerator	Appcelerator Studio	AS	2919	1	40	5.64	5	3.33
	Aptana Studio	AP	829	1	40	8.02	8	5.95
	Titanium SDK/CLI	TI	2251	1	34	6.32	5	5.10
DuraSpace	DuraCloud	DC	666	1	16	2.13	1	2.03
Atlassian	Bamboo	BB	521	1	20	2.42	2	2.14
	Clover	CV	384	1	40	4.59	2	6.55
	JIRA Software	JI	352	1	20	4.43	3	3.51
Moodle	Moodle	MD	1166	1	100	15.54	8	21.65
Lsstcorp	Data Management	DM	4667	1	100	9.57	4	16.61
Mulesoft	Mule	MU	889	1	21	5.08	5	3.50
	Mule Studio	MS	732	1	34	6.40	5	5.39
Spring	Spring XD	XD	3526	1	40	3.70	3	3.23
Talendforge	Talend Data Quality	TD	1381	1	40	5.92	5	5.19
	Talend ESB	TE	868	1	13	2.16	2	1.50
Total			23,313					

16 projects highlights the different levels of task complexity and the varying nature of project management within each repository.

4. Proposed Methodology

In this work, we present SPERT, a model designed to automatically estimate story points by processing raw textual data from user story titles and descriptions. SPERT integrates state-of-the-art transformer embeddings with adaptive regression based on RL, allowing the model to dynamically adjust to evolving project environments and make accurate predictions in real-time.

The architecture of SPERT is illustrated in Fig. 3, which highlights its core components: pretrained transformer-based embeddings, a multi-layer transformer encoder and an adaptive regression layer enhanced by RL.

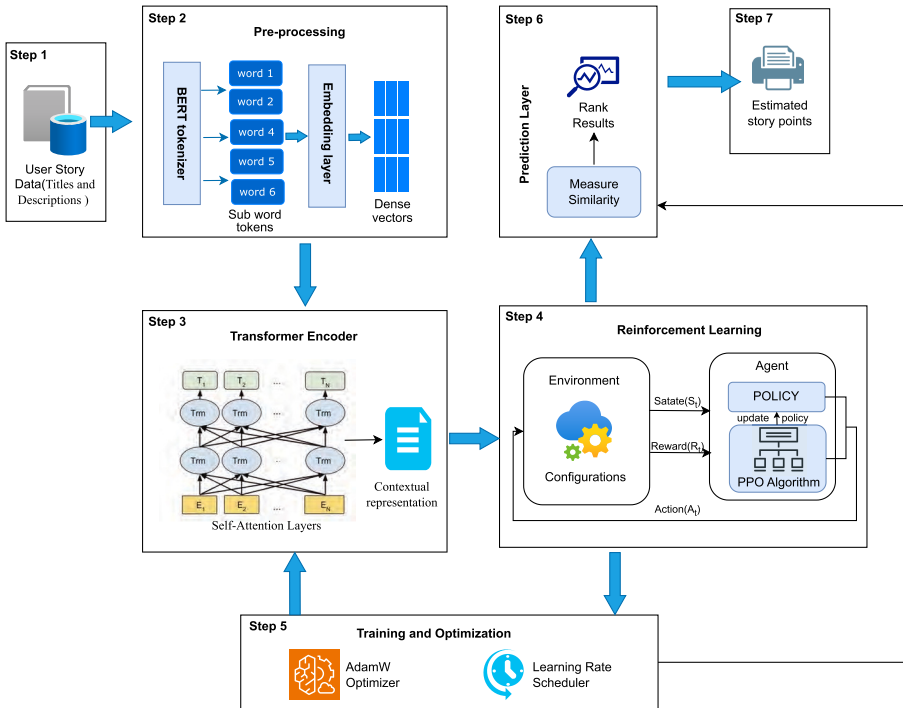


Fig. 3. Architecture of the SPERT model, highlighting transformer embeddings, the multi-layer transformer encoder and the RL-based adaptive regression layer.

4.1. Pretrained transformer embeddings

SPERT leverages pretrained transformer embeddings from BERT, which is designed to capture deep semantic representations of text. These transformers are pretrained on vast corpora, allowing them to learn rich contextual information, such as

syntactic relationships, word meanings and the broader context within user stories. This ensures that the embeddings can represent both simple and complex relationships between words in user stories.

The input to SPERT consists of the concatenated title and description of each user story, formatted with BERT’s special tokens as follows:

$$[\text{CLS}]T[\text{SEP}]D[\text{SEP}],$$

where

- **[CLS]**: A classification token added at the beginning of the sequence. Its embedding is used by SPERT to represent the overall sequence context for story point estimation.
- **[SEP]**: A separator token used to distinguish the title and description. It also helps BERT understand the boundaries between segments.

The sequence is tokenized and padded to a fixed length of 128 tokens, ensuring uniform input sizes for BERT. The resulting embedding matrix captures both semantic and syntactic relationships, as detailed below.

The formal embedding process is represented as

$$E = \text{Embedding}([\text{CLS}] \oplus T \oplus [\text{SEP}] \oplus D \oplus [\text{SEP}]), \quad (1)$$

where E is the resulting embedding, and T and D represent the title and description of the user story, respectively. These embeddings are fed into the transformer encoder, which extracts further features relevant for the story point estimation task.

Token Limit Justification: The input sequence length for SPERT is set to a maximum of 128 tokens, which includes BERT’s special tokens ([CLS] and [SEP]). This limit ensures computational efficiency while capturing sufficient information for the task. Analysis of the dataset revealed that 95% of user stories (titles and descriptions combined) are represented within 128 tokens, including special tokens.

In cases where the content exceeds 128 tokens, the text is truncated, prioritizing the initial portion of the sequence, as it often contains the most relevant information. To validate the sufficiency of this limit, we conducted experiments using larger token limits (256 and 512 tokens). These experiments showed no significant improvement in prediction accuracy, confirming that a 128-token limit effectively balances information retention and computational efficiency.

Role of Syntactic Information: Although BERT is primarily designed to capture semantic information, it implicitly encodes syntactic relationships through its self-attention mechanism and pretraining objectives. For example, the transformer architecture processes the relative positions and contextual dependencies of words, which inherently reflect syntactic patterns. The use of special tokens such as ‘[CLS]’ and ‘[SEP]’ further supports this by marking segment boundaries and providing positional cues, aiding in understanding the structure and organization of text.

In the context of SPERT, syntactic information plays a complementary role in enhancing semantic understanding. User stories often describe tasks with specific structures and action-oriented phrasing, where syntactic cues such as the order of verbs, modifiers and objects can hint at task complexity. For instance, a sentence like “optimize database performance by reducing query times” contains syntactic patterns that emphasize the relationship between actions (optimize, reducing) and objects (database, query times), indirectly signaling effort and complexity.

By leveraging BERT embeddings, SPERT benefits from a rich representation of both syntactic and semantic features, even though syntax is not the primary focus of the model. This dual understanding enables SPERT to better analyze user stories and produce accurate story point predictions.

4.2. Transformer encoder for document representation

Once the embeddings are generated, they are passed through a multi-layer transformer encoder. This encoder consists of multiple layers of self-attention mechanisms and feedforward networks. The self-attention mechanism is particularly powerful in capturing long-range dependencies within the text, which is important for understanding the nuances of user stories in Agile development.

The self-attention mechanism assigns different weights to various parts of the input text, allowing the model to focus on the most relevant sections when making predictions. For example, in a user story, the phrase “implement user login” may carry more weight than descriptive context such as “in the next sprint”.

The output of the transformer encoder is denoted as

$$H = \text{TransformerEncoder}(E). \quad (2)$$

This output encodes both the semantic content of the input and the structural relationships between words, making it a rich representation for further processing. The multi-layer transformer encoder’s depth ensures that even complex dependencies between words are captured effectively, enhancing the model’s ability to make accurate story point predictions.

4.3. Enhanced explanation of BERT and RL integration in SPERT

The integration of BERT embeddings and RL in SPERT forms the core of its adaptability and prediction accuracy. This section provides a detailed explanation of their interaction and a flowchart to illustrate the information flow within the model.

BERT Embedding Generation: SPERT utilizes pretrained BERT embeddings to capture the semantic richness and context of user stories. The input to the model, consisting of the title and description of a user story, is tokenized using BERT’s tokenizer. This tokenized input is passed through the BERT encoder, producing dense vector embeddings that represent both the syntactic and semantic features of the user story. These embeddings are denoted as $E = \text{BERT}(T \oplus D)$, where T represents the title, D represents the description and \oplus denotes concatenation.

Interaction with RL: The BERT embeddings (E) serve as the primary input to the RL-based adaptive regression layer. The RL component, implemented using PPO, dynamically adjusts predictions based on feedback from the Mean Absolute Error (MAE). The interaction proceeds as follows:

- (1) **Initial Prediction:** The embeddings generated by BERT are processed through a series of fully connected layers to produce an initial story point estimate (\hat{y}_{init}).
- (2) **Feedback Mechanism:** The RL agent evaluates the initial prediction by comparing it to the ground truth (y) using the MAE loss. This feedback signal guides the adjustment of the prediction.
- (3) **Policy Update:** PPO updates the policy network to minimize prediction errors. This involves refining the weight parameters such that future predictions align more closely with the actual story points.
- (4) **Final Prediction:** The updated policy generates a refined story point estimate (\hat{y}_{final}), ensuring that the model remains adaptable to dynamic project conditions.

The RL mechanism operates as a dynamic layer that learns from project-specific feedback, allowing SPERT to continuously improve its predictions during training.

Figure 3 illustrates the overall architecture of SPERT, encompassing all major components and their interactions. Steps 3 and 4 in the architecture emphasize the critical relationship between the BERT embedding layer and the RL-based regression layer. Specifically, the embeddings generated by BERT serve as input to the RL mechanism, which dynamically refines predictions based on feedback. This interplay allows SPERT to continuously adapt and improve its story point estimations.

4.4. Adaptive regression with RL

SPERT's adaptive regression layer employs RL to dynamically adjust the model's predictions based on feedback. This enables SPERT to handle evolving project environments that are typical in Agile development, where user story complexities can change frequently.

SPERT uses the PPO algorithm for RL. PPO is chosen because of its stability and efficiency in learning policies for continuous control tasks [45], which in our case translates to learning how to adjust story point estimates over time.

How SPERT Adapts:

- (1) **Initial Prediction:** The transformer encoder produces an initial story point estimate based on the embeddings and document representation.
- (2) **Feedback Mechanism:** The PPO algorithm evaluates the prediction by comparing it to the actual story point and computing the MAE. This error serves as the feedback signal.

- (3) **Policy Update:** Based on the feedback, the PPO algorithm updates the model's policy to improve future predictions. The objective is to minimize prediction errors, and PPO optimizes the following objective function:

$$L^{\text{PPO}} = \mathbb{E}_t[\min(r_t(\theta)A_t, \text{clip}(r_t(\theta), 1 - \epsilon, 1 + \epsilon)A_t)]. \quad (3)$$

- (4) **Continuous Adaptation:** As new user stories and project conditions emerge, the RL component refines the predictions, making the model highly adaptable. This is particularly important in Agile environments, where the nature of user stories and project scopes can shift dramatically over time.

Structuring Feedback as Rewards or Penalties: In the RL framework, project feedback is structured into a reward signal that guides the policy updates. The reward function is defined to reflect the prediction accuracy, with higher rewards assigned to predictions closer to the actual story points and penalties applied to larger deviations. Specifically, the reward R is formulated as

$$R = -|y - \hat{y}|, \quad (4)$$

where y is the actual story point and \hat{y} is the predicted story point. This formulation ensures that the RL agent is incentivized to minimize the absolute error.

Balancing Prior Knowledge and New Feedback: SPERT achieves a balance between learning from prior predictions and adapting to new feedback through the PPO algorithm. The policy update process incorporates both the cumulative reward from past predictions and the instantaneous reward from new project feedback. By maintaining a buffer of historical experiences, SPERT prevents overfitting to recent data while still adapting to evolving project conditions. The clipped objective function in PPO ensures stable updates, avoiding drastic policy changes that could destabilize the learning process.

This balance allows SPERT to generalize effectively across projects while retaining the flexibility to refine its predictions dynamically as new feedback becomes available.

By integrating this structured feedback loop, SPERT not only refines its predictions over time but also ensures that the learning process remains robust and adaptable to the dynamic nature of Agile environments.

4.5. Final prediction layer

After the document representation has been refined through RL, it is passed through an Actor-Critic model, which processes the final representation, R , and outputs the continuous value representing the estimated story point:

$$y = \text{ActorCritic}(R). \quad (5)$$

This final output accounts for the inherent complexity and variability of tasks in Agile development, providing a robust estimate of the story points.

4.6. System training and optimization

The SPERT model is trained using the AdamW optimizer, which is specifically suited for transformer-based models. AdamW combines the adaptive learning rate strategy of Adam with weight decay, preventing overfitting during training.

A learning rate scheduler is used, gradually adjusting the learning rate during training to ensure smooth convergence. Dropout regularization is also employed to avoid overfitting, particularly when dealing with smaller datasets or when the model is prone to memorizing specific patterns.

The training objective is to minimize the MAE:

$$\mathcal{L}(\hat{y}, y) = \frac{1}{n} \sum_{i=1}^n |\hat{y}_i - y_i|. \quad (6)$$

Through RL, the model continues to adapt during training, improving performance over time and making it highly suitable for dynamic project environments.

5. Experimental Setup

In this section, we describe the experimental framework used to train, evaluate and validate SPERT. The setup ensures robustness and scalability across a diverse dataset of Agile project stories and leverages state-of-the-art techniques in NLP and RL.

5.1. Preprocessing

The raw textual data are processed to ensure that it is clean and suitable for transformer-based models. Preprocessing includes standardizing casing, removing special characters and truncating long sequences. Key steps are the following:

- **Text Cleaning:** Removal of HTML tags, URLs and non-alphanumeric characters to reduce noise.
- **Tokenization:** Tokenization is performed using the BERT tokenizer, converting text into token sequences for transformer-based processing.
- **Padding and Truncation:** Sequences are either padded or truncated to a fixed length of 128 tokens to ensure uniform input sizes.

These preprocessing steps ensure that the textual data are clean, standardized and ready for feature extraction by the transformer models [18].

5.2. Data splitting for cross-project generalization

The dataset was split into training, validation and test sets using a 60-20-20 ratio. To avoid data leakage, we used a project-level split, meaning that entire projects were kept in separate splits. This simulates a real-world scenario, where the model is evaluated on new projects it has not encountered during training.

The split was stratified by project type and domain to maintain a balanced representation across the different Agile projects, allowing the model to generalize its learning across diverse domains.

5.3. Feature extraction with transformer embeddings

SPERT utilizes pretrained BERT transformer [11, 19], for feature extraction. These models provide contextual embeddings that capture both syntactic and semantic relationships in the user stories. Transformers are particularly well-suited for this task due to their ability to handle long-range dependencies, which are crucial for understanding complex stories in Agile environments.

By leveraging pretrained embeddings, SPERT avoids the need for manual feature engineering and ensures scalability across diverse projects.

5.4. Hyperparameter tuning

To optimize the model's performance, a grid search was conducted to explore various hyperparameter combinations, such as batch size, learning rate and dropout rate. The optimal configuration was selected based on the model's performance on the validation set. The key hyperparameters tested during the grid search are shown below.

Hyperparameter Tuning Grid Search:

- **Learning Rate:** $\{1e-5, 3e-5, 5e-5\}$
- **Batch Size:** $\{16, 32\}$
- **Dropout Rate:** $\{0.1, 0.3\}$

Basis for Adjustment Range and Step Size: The adjustment ranges and step sizes for these hyperparameters were determined based on a combination of prior research, preliminary experiments and practical considerations:

- **Learning Rate:** The range of $\{1e-5, 3e-5, 5e-5\}$ was selected based on empirical evidence from related works involving transformer models such as BERT [18]. These values are known to provide stable convergence during fine-tuning. The step size of $2e-5$ was chosen to ensure adequate coverage of the range without unnecessary computational overhead.
- **Batch Size:** The range of $\{16, 32\}$ reflects common choices for BERT-based models, where computational efficiency and gradient stability are critical. Larger batch sizes were avoided to prevent memory issues, while the step size of 16 aligns with practical limitations of GPU hardware used in the experiments.
- **Dropout Rate:** The values $\{0.1, 0.3\}$ were selected based on their effectiveness in regularizing deep models without causing underfitting or overfitting [46]. The step size of 0.2 was considered sufficient to test low and moderate dropout settings while keeping the grid search computationally manageable.

These ranges and step sizes ensure a balance between exhaustive search and computational feasibility, allowing the model to generalize effectively while avoiding overfitting of hyperparameters. The final configuration was selected based on the combination yielding the lowest validation error during the grid search.

5.5. Training configuration

The model was trained using the AdamW optimizer [47, 48], with a learning rate scheduler and warm-up steps to ensure smooth convergence. Dropout regularization was applied to prevent overfitting, and the overall training objective was to minimize MAE. The training hyperparameters are summarized below.

Training Hyperparameters:

- **Learning Rate:** $3e-5$
- **Batch Size:** 32
- **Dropout Rate:** 0.1
- **Epochs:** 20

6. Model Training

This section outlines the training process of SPERT, covering key components such as training objectives, optimization strategies and the integration of RL. The goal of the training process is to ensure that the model learns effectively from the data while avoiding overfitting or underfitting.

6.1. Training objectives

The primary goal of the training procedure is to minimize the error in story point estimation. MAE is used as the loss function:

$$\mathcal{L}(\hat{y}, y) = \frac{1}{n} \sum_{i=1}^n |\hat{y}_i - y_i|. \quad (7)$$

MAE is preferred because it provides a more interpretable measure of error for tasks like story point estimation. Agile projects often have varying degrees of task sizes and complexities, and MAE captures these variances more effectively than Mean Squared Error (MSE), which can overemphasize outliers.

6.2. Training procedure

SPERT's training process is divided into two stages: transformer-based training for feature extraction and RL-based adaptive regression. This two-stage approach allows the model to capture semantic relationships and adjust predictions dynamically based on feedback.

6.2.1. Stage 1: Transformer-based training

In the first stage, the transformer embeddings are fine-tuned using the AdamW optimizer. The transformer layers are updated based on the training data, capturing essential features for story point estimation.

Hyperparameters for Fine-tuning:

- **Batch Size:** 32
- **Learning Rate:** $3e-5$
- **Dropout Rate:** 0.1
- **Epochs:** 20

During this stage, dropout and weight decay are used to prevent overfitting. The fine-tuned embeddings provide a rich representation of the user stories for the regression layer.

6.2.2. Fine-tuning BERT results across projects

To evaluate the effectiveness of the BERT fine-tuning process across all 16 projects, we monitored the training loss for 20 epochs. The combined loss curves for all projects are presented in Fig. 4. This figure demonstrates consistent reduction in training loss, with most projects converging by epoch 15.

Upon analysis of the loss curves, it was observed that the loss stopped decreasing significantly after approximately 15 epochs for most projects. This indicates that training beyond 15 epochs offers diminishing returns, making it a suitable stopping

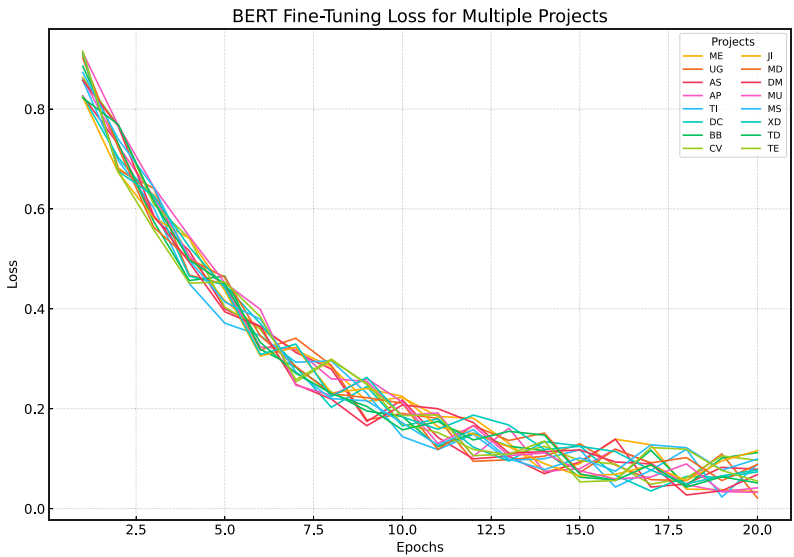


Fig. 4. Combined fine-tuning loss curves for all 16 projects during the BERT fine-tuning process.

point for fine-tuning. The results highlight the robustness of the fine-tuning process, showing that SPERT effectively optimized the BERT embeddings for semantic representation across diverse Agile project datasets.

6.2.3. Stage 2: RL-based adaptive regression

Once the transformer embeddings are trained, the adaptive regression layer is trained using RL. The PPO algorithm is used to update the policy based on feedback from the MAE loss.

The PPO algorithm optimizes the policy π by maximizing the following objective:

$$L^{\text{PPO}} = \mathbb{E}_t[\min(r_t(\theta)A_t, \text{clip}(r_t(\theta), 1 - \epsilon, 1 + \epsilon)A_t)]. \quad (8)$$

This allows SPERT to adapt its predictions over time, ensuring that the model remains responsive to changing project conditions.

6.3. Optimization strategy

The AdamW optimizer is used for both the transformer fine-tuning and RL stages. The learning rate scheduler is applied to ensure smooth convergence, with the following schedule:

$$\eta_t = \eta_0 \cdot \left(1 - \frac{t}{T}\right). \quad (9)$$

This strategy allows the model to start with a smaller learning rate and gradually decay it over time to avoid overfitting.

6.4. Training time and computational resources

The training process was conducted on NVIDIA Tesla V100 GPUs, with a total training time of 30 h. Training was distributed across four GPUs to reduce computational overhead and ensure timely convergence.

Training Time:

- **Transformer Fine-tuning Stage:** 10 h
- **RL Stage:** 20 h

7. Evaluation

To evaluate how accurate the SPERT predictions are, we use several evaluation metrics: MAE, Median Absolute Error (MdAE) and SA, as shown in the following formulas. These metrics are chosen because they provide an unbiased and reliable way to evaluate prediction models, even in situations involving under- or over-estimations, which is essential for a fair assessment process as supported by previous SEE studies [32, 49–51].

- MAE measures the average magnitude of the errors in the story point predictions, without considering their direction. It is defined as

$$\text{MAE} = \frac{1}{n} \sum_{i=1}^n |y_i - \hat{y}_i|, \quad (10)$$

where y_i is the actual value, \hat{y}_i is the predicted value and n is the number of observations. MAE provides an intuitive measure of model accuracy, with lower values indicating better performance.

- MdAE is a robust measure of prediction error, which is less sensitive to outliers compared to MAE. It is calculated as the median of the absolute differences between actual and predicted values:

$$\text{MdAE} = \text{median}(|y_1 - \hat{y}_1|, |y_2 - \hat{y}_2|, \dots, |y_n - \hat{y}_n|). \quad (11)$$

MdAE provides a better representation of typical prediction errors, especially when the dataset contains extreme values.

- SA is a metric that reflects the percentage of predictions that fall within an acceptable range of the actual values, making it a practical measure for assessing estimation accuracy. It is defined as

$$\text{SA} = \frac{1}{n} \sum_{i=1}^n \mathbb{I}(|y_i - \hat{y}_i| < \delta), \quad (12)$$

where \mathbb{I} is an indicator function that returns 1 if the condition is met, and 0 otherwise, and δ is a predefined threshold that represents an acceptable error margin.

RQ1: How does SPERT compare with state-of-the-art ML and NLP-based methods for story point estimation in Agile projects?

To address this question, we conducted a comprehensive comparative evaluation of SPERT against several baseline models, including SBERT-XG, LHC-SE, Deep-SE and TF-IDF-SE. The primary objective of this evaluation was to determine how well SPERT performs in terms of story point estimation accuracy compared to existing state-of-the-art methods.

The evaluation experiments were carried out across multiple Agile projects to assess the generalizability of SPERT's performance. Each of the baseline models was trained and tested under the same conditions as SPERT to ensure fairness in comparisons. The experiments involved training the models on 60% of the dataset, while the remaining 20% was used for testing and 20% for validation. We further conducted cross-project evaluations to assess the ability of SPERT to generalize to new and unseen projects.

In addition, we conducted statistical significance testing using the Wilcoxon signed-rank test to determine whether SPERT's performance improvements were statistically significant compared to the baseline models. We also used the A12 effect

size to measure the practical relevance of the observed improvements, with values greater than 0.7 considered indicative of substantial practical gains.

RQ2: Does the use of RL enhance the predictive capability of SPERT compared to traditional regression approaches?

To answer this question, we compared SPERT with its RL component enabled and disabled. Specifically, we evaluated the model's performance with and without the use of PPO in the adaptive regression layer. This allowed us to isolate the impact of RL on prediction accuracy and adaptability.

The comparative evaluation was conducted using the same performance metrics as for RQ1, i.e. MAE, MdAE and SA. We trained both versions of the model on the same dataset and performed cross-validation to assess their predictive performance. We then measured the difference in accuracy between the two versions, focusing on whether the RL-enhanced SPERT was able to achieve better story point estimates by learning from dynamic feedback.

RQ3: How does SPERT's transformer-based embeddings compare to traditional NLP techniques like TF-IDF and Doc2Vec for story point estimation?

To evaluate the effectiveness of SPERT's transformer-based embeddings, we compared it against traditional NLP techniques, such as TF-IDF and Doc2Vec. The key focus of this evaluation was to determine whether the contextual embeddings generated by transformer models like BERT provided a more accurate representation of user stories compared to the traditional, static representations.

We ran experiments with the same evaluation metrics (MAE, MdAE and SA) and measured the performance differences between SPERT, TF-IDF-SE and SBERT-XG across multiple projects. These experiments were designed to highlight the benefits of deep contextual understanding provided by transformers in predicting the complexity of Agile user stories.

RQ4: How well does SPERT generalize across multiple Agile projects compared to baseline models, especially in cross-project estimation scenarios?

To assess SPERT's cross-project generalization capabilities, we conducted cross-project estimation experiments. In these experiments, SPERT was trained on one project and tested on another to evaluate its ability to generalize beyond the specific dataset it was trained on.

We compared SPERT's performance to the baseline models (SBERT-XG, LHC-SE, Deep-SE and TF-IDF-SE) to evaluate its robustness. The cross-project evaluations provide insights into how well the model can adapt to new and unseen Agile projects, which is crucial for its applicability in real-world Agile software development scenarios.

The metrics used for these cross-project experiments included MAE and MdAE, while we also conducted the Wilcoxon signed-rank test and computed the A12 effect size to assess the statistical and practical significance of SPERT's generalization capabilities.

8. Results

In this section, we present the detailed results of the experiments described in the previous section. We provide a thorough analysis of SPERT's performance in comparison with the baseline models and discuss the implications of the findings in terms of story point estimation accuracy, adaptability and generalizability.

RQ1: SPERT comparison with state-of-the-art ML and NLP-based methods

The comparison between SPERT and the baseline models (SBERT-XG, LHC-SE, Deep-SE and TF-IDF-SE) is summarized in Table 2. SPERT consistently demonstrated superior performance across all evaluated metrics. Specifically, SPERT achieved a MAE of 0.71 on the Mesos project, compared to 1.64 for SBERT-XG, 1.34 for LHC-SE and 1.02 for Deep-SE.

The MdAE and SA metrics further indicate that SPERT provides more reliable story point predictions. For instance, SPERT's MdAE was significantly lower than the baselines, demonstrating that the model is less sensitive to extreme outliers.

SPERT demonstrates consistent improvements over baseline models across multiple projects due to several factors. First, the use of pretrained BERT embeddings enables SPERT to capture both semantic and syntactic relationships within user stories, providing a richer understanding of the text compared to traditional techniques like TF-IDF or shallow embeddings. Additionally, SPERT's RL mechanism dynamically adjusts predictions based on feedback, allowing the model to generalize better across diverse project environments. Finally, the inclusion of advanced preprocessing steps, such as tokenization and truncation, ensures clean and uniform inputs, further improving performance.

The degree of improvement varies across projects due to the diversity in datasets. Projects like ME and UG have user stories that align closely with the linguistic patterns learned by BERT during pretraining, resulting in significant performance gains. Conversely, projects with specialized or domain-specific language, such as MD and DM, may benefit less from BERT's general-purpose embeddings. Additionally, the variability in the distribution of story points across projects can affect SPERT's ability to generalize, as projects with balanced story point distributions provide more robust training signals.

SPERT's Performance on the DM Project: While SPERT generally outperforms baseline models, its performance on the DM project was slightly worse. Analysis of the DM dataset revealed a high concentration of outlier story points, which introduced noise into the training process and adversely affected the model's predictions. Additionally, the DM project involves domain-specific terminology and unique patterns that were less effectively captured by BERT embeddings. In contrast, baseline models like SBERT-XG, which utilize simpler algorithms, were less sensitive to outliers, resulting in slightly better performance. This highlights the importance of fine-tuning SPERT further for projects with unique characteristics or developing strategies to mitigate the impact of outliers.

Table 2. Overall performance of SPERT and baseline models.

Project	Abb.	Model	MAE	MdAE	SA	Project	Abb.	Model	MAE	MdAE	SA
Mesos	ME	SPERT	0.71	0.49	71.22	JIRA Software	JI	SPERT	0.98	0.76	65.36
		SBERT-XG	1.64	1.25	63.03			SBERT-XG	1.36	1.45	56.32
		LHC-SE	1.34	1.00	34.38			LHC-SE	1.58	1.98	49.57
		Deep-SE	1.02	0.73	59.84			Deep-SE	1.38	1.09	59.52
		TF-IDF-SE	1.34	1.00	34.38			TF-IDF-SE	1.47	1.62	46.12
		Mean	1.64	1.78	35.61			Mean	2.48	2.15	27.06
Usergrid	UG	Median	1.73	2.00	32.01	Moodle	MD	Median	2.93	2.00	13.88
		SPERT	0.68	0.46	67.59			SPERT	6.04	2.98	68.14
		SBERT-XG	1.65	1.34	48.96			SBERT-XG	7.21	3.25	63.1
		LHC-SE	2.05	1.97	44.65			LHC-SE	6.31	7.00	57.3
		Deep-SE	1.03	0.80	52.66			Deep-SE	5.97	4.93	50.29
		TF-IDF-SE	1.65	1.42	46.36			TF-IDF-SE	6.31	7.00	57.3
Appcelerator Studio	AS	Mean	1.48	1.23	32.13	Data Management	DM	Mean	10.9	12.11	9.16
		Median	1.60	1.00	26.29			Median	7.18	6.00	40.16
		SPERT	1.24	1.02	67.54			SPERT	1.77	1.52	81.23
		SBERT-XG	1.71	1.29	54.62			SBERT-XG	2.79	2.06	74.51
		LHC-SE	1.51	2.00	51.89			LHC-SE	1.56	1.00	53.87
		Deep-SE	1.36	0.58	60.26			Deep-SE	3.77	2.22	47.87
Aptana Studio	AP	TF-IDF-SE	1.51	2.00	51.89	Mule	MU	TF-IDF-SE	1.49	1.00	55.71
		Mean	2.08	1.52	39.02			Mean	5.29	4.55	26.85
		Median	1.84	1.00	46.17			Median	4.82	3.00	33.38
		SPERT	1.19	0.11	76.73			SPERT	1.89	1.47	56.74
		SBERT-XG	5.81	4.49	63.86			SBERT-XG	2.67	2.37	42.00
		LHC-SE	4.14	3.00	3.20			LHC-SE	2.27	2.00	37.11
		Deep-SE	2.71	2.52	42.58			Deep-SE	2.18	1.96	40.09
		TF-IDF-SE	3.99	3.00	32.81			TF-IDF-SE	3.58	2.00	0.81
		Mean	3.15	3.46	33.3			Mean	2.59	2.22	28.82
		Median	3.71	4.00	21.54			Median	2.69	2.00	26.07

Table 2. (Continued)

Project	SDK/CLI	Abb.	Model	MAE	MdAE	SA	Project	Abb.	Model	MAE	MdAE	SA
Titanium	CLI	TI	SPERT	1.37	1.02	78.32	Mule Studio	MS	SPERT	2.58	2.14	68.83
			SBERT-XG	2.1	1.57	51.53			SBERT-XG	3.56	3.21	33.20
			LHC-SE	2.53	2.00	30.70			LHC-SE	3.62	3.41	30.02
			Deep-SE	1.97	1.34	55.92			Deep-SE	3.23	1.99	17.17
			TF-IDF-SE	2.53	2.00	30.7			TF-IDF-SE	3.58	3.47	28.41
			Mean	3.05	1.97	31.59			Mean	3.34	2.68	14.21
DuraCloud		DC	Median	2.47	2.00	44.65	Spring XD	XD	Median	3.30	2.98	15.42
			SPERT	1.15	0.64	75.08			SPERT	1.38	1.02	71.61
			SBERT-XG	2.15	1.90	19.61			SBERT-XG	1.91	1.57	67.23
			LHC-SE	1.25	1.00	9.65			LHC-SE	1.54	1.00	39.53
			Deep-SE	0.68	0.70	69.92			Deep-SE	1.63	1.31	46.82
			TF-IDF-SE	0.68	1.00	39.94			TF-IDF-SE	2.01	2.00	20.82
Bamboo		BB	Mean	1.30	1.14	42.88	Talend Data Quality	TD	Mean	2.27	2.53	26.00
			Median	0.73	1.00	68.08			Median	2.07	2.00	32.55
			SPERT	0.46	0.41	71.25			SPERT	2.41	2.18	69.23
			SBERT-XG	1.23	1.45	56.32			SBERT-XG	3.63	2.31	48.83
			LHC-SE	1.14	1.56	52.02			LHC-SE	3.52	3.00	27.6
			Deep-SE	0.74	0.61	71.24			Deep-SE	2.97	2.92	48.28
Clover		CV	TF-IDF-SE	1.02	1.13	63.84	Talend ESB	TE	TF-IDF-SE	5.05	5.00	3.95
			Mean	1.75	1.31	32.11			Mean	4.81	5.00	16.18
			Median	1.32	1.00	48.72			Median	3.87	4.00	32.43
			SPERT	1.68	1.23	81.24			SPERT	0.36	0.20	88.71
			SBERT-XG	2.51	1.69	75.21			SBERT-XG	1.91	1.72	63.31
			LHC-SE	3.88	2.00	46.35			LHC-SE	0.99	1.00	32.56
			Deep-SE	2.11	0.8	50.45			Deep-SE	0.64	0.59	69.67
			TF-IDF-SE	4.04	1.00	44.15			TF-IDF-SE	0.97	1.00	33.95
			Mean	3.49	3.06	17.84			Mean	1.14	0.91	45.86
			Median	2.84	2.00	33.33			Median	1.16	1.00	44.44

The statistical significance of these results was confirmed using the Wilcoxon signed-rank test, with p -values below 0.05 for most comparisons, indicating that SPERT's performance improvements were statistically significant. The A12 effect size values, averaging around 0.72, further confirm the practical significance of these improvements.

Figure 5 illustrates the overall performance comparison in terms of MAE between SPERT and the baseline models.

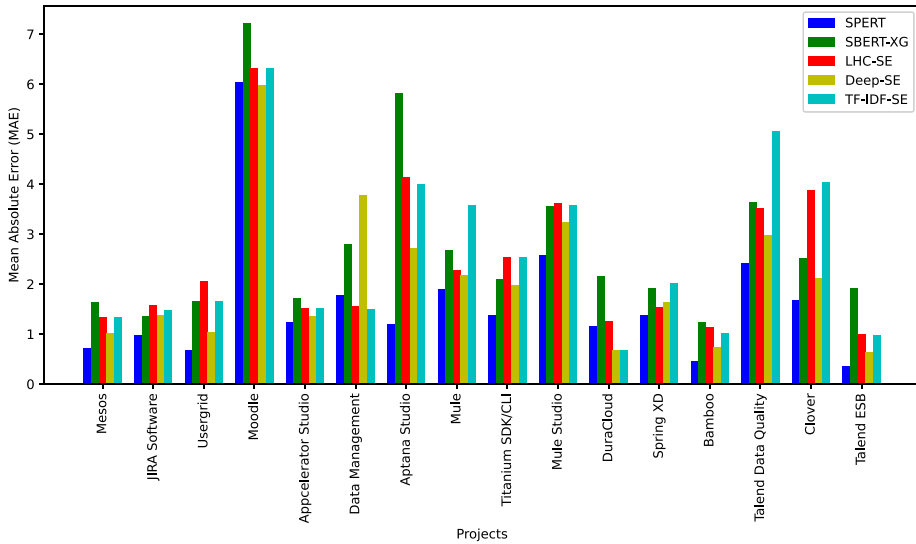


Fig. 5. Overall performance comparison (MAE) between SPERT and baseline models.

RQ2: Use of RL enhances the predictive capability of SPERT

The result of the RL component on SPERT's performance is shown in Table 3. The results indicate that SPERT with RL consistently outperformed the version without RL across all projects. For instance, SPERT with RL achieved an MAE of 0.68 in the Usergrid project, compared to 0.95 without RL, showing a significant improvement in prediction accuracy.

Impact of RL Across Projects: The impact of RL varies significantly across projects due to differences in the characteristics of the datasets and feedback signals. Projects with highly structured user stories, such as ME and UG, benefit more from RL as it dynamically refines predictions by leveraging project-specific feedback. In contrast, projects like DM and MD, which contain outlier story points or domain-specific language, see smaller improvements with RL due to the increased complexity of learning stable policies in noisy environments. Additionally, the variability in reward quality — affected by the accuracy of initial predictions and the distribution of story points — can influence the degree to which RL contributes to performance gains.

Table 3. Impact of RL on model performance.

Proj	w/o RL	w/RL	Improvement
	MAE/MdAE/SA	MAE/MdAE/SA	MAE/MdAE/SA
ME	0.89/0.70/60.1	0.71/0.49/71.2	-0.18/-0.21/+11.1
UG	0.95/0.75/60.0	0.68/0.46/67.6	-0.27/-0.29/+7.6
AS	1.50/1.20/63.4	1.24/1.02/67.5	-0.26/-0.18/+4.1
AP	1.50/1.20/65.5	1.19/0.11/76.7	-0.31/-1.09/+11.2
TI	1.67/1.30/65.5	1.37/1.02/78.3	-0.30/-0.28/+12.8
DC	1.35/0.89/62.5	1.15/0.64/75.1	-0.20/-0.25/+12.6
BB	0.60/0.51/61.0	0.46/0.41/71.3	-0.14/-0.10/+10.3
CV	2.00/1.60/75.5	1.68/1.23/81.2	-0.32/-0.37/+5.7
JI	1.20/1.00/60.5	0.98/0.76/65.4	-0.22/-0.24/+4.9
MD	6.20/3.40/65.5	6.04/2.98/68.1	-0.16/-0.42/+2.6
DM	2.00/1.70/75.0	1.77/1.52/81.2	-0.23/-0.18/+6.2
MU	2.20/1.70/52.0	1.89/1.47/56.7	-0.31/-0.23/+4.7
MS	2.80/2.30/63.0	2.58/2.14/68.8	-0.22/-0.16/+5.8
XD	1.50/1.20/65.5	1.38/1.02/71.6	-0.12/-0.18/+6.1
TD	2.80/2.50/63.0	2.41/2.18/69.2	-0.39/-0.32/+6.2
TE	0.45/0.31/75.0	0.36/0.20/88.7	-0.09/-0.11/+13.7

Comparison of Fine-Tuned BERT Without RL: Fine-tuned BERT alone demonstrates significant improvements over baseline models due to its ability to capture rich semantic and syntactic relationships within user stories. This raises an important consideration for resource-constrained environments: while RL offers additional benefits, especially in projects requiring dynamic adaptation to feedback, fine-tuned BERT can achieve competitive results at a lower computational cost. For example, in projects like BB and TI, fine-tuned BERT achieves performance comparable to SPERT with RL, highlighting its sufficiency for projects with stable or less complex datasets.

Recommendations for Constrained Training Conditions: In scenarios with limited computational resources, directly fine-tuning BERT may be a viable alternative to SPERT with RL. This approach retains the majority of performance gains while significantly reducing training complexity and resource demands. However, for projects with highly dynamic or noisy environments, RL remains essential for handling evolving conditions and improving generalization.

The MAE improvement is visualized in Fig. 6, showing that the RL-enhanced SPERT has lower error rates compared to its non-RL counterpart across different projects.

RQ3: SPERT comparison with traditional NLP techniques

The comparative evaluation of SPERT’s transformer-based embeddings with traditional NLP techniques (TF-IDF and Doc2Vec) is summarized in Table 4. SPERT outperformed both TF-IDF-SE and Doc2Vec in all evaluated metrics. For instance, SPERT achieved an MAE of 1.24 in the Appcelerator Studio project, compared to 1.51 for TF-IDF-SE and 1.71 for SBERT-XG.

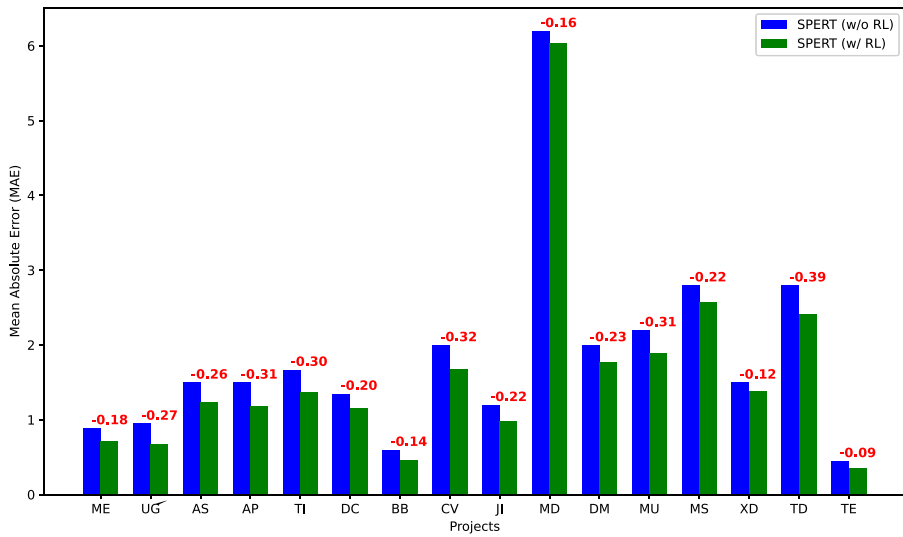


Fig. 6. Impact of RL on model performance (MAE).

Table 4. Performance comparison: Transformer-based models vs traditional NLP models.

Abb.	Model	MAE	MdAE	SA	Abb.	Model	MAE	MdAE	SA
ME	SPERT	0.71	0.49	71.22	JI	SPERT	0.98	0.76	65.36
	SBERT-XG	1.64	1.25	63.03		SBERT-XG	1.36	1.45	56.32
	TF-IDF-SE	1.34	1.00	34.38		TF-IDF-SE	1.47	1.62	46.12
UG	SPERT	0.68	0.46	67.59	MD	SPERT	6.04	2.98	68.14
	SBERT-XG	1.65	1.34	48.96		SBERT-XG	7.21	3.25	63.10
	TF-IDF-SE	1.65	1.42	46.36		TF-IDF-SE	6.31	7.00	57.30
AS	SPERT	1.24	1.02	67.54	DM	SPERT	1.77	1.52	81.23
	SBERT-XG	1.71	1.29	54.62		SBERT-XG	2.79	2.06	74.51
	TF-IDF-SE	1.51	2.00	51.89		TF-IDF-SE	1.49	1.00	55.71
AP	SPERT	1.19	0.11	76.73	MU	SPERT	1.89	1.47	56.74
	SBERT-XG	5.81	4.49	63.86		SBERT-XG	2.67	2.37	42.00
	TF-IDF-SE	3.99	3.00	32.81		TF-IDF-SE	3.58	2.00	0.81
TI	SPERT	1.37	1.02	78.32	MS	SPERT	2.58	2.14	68.83
	SBERT-XG	2.10	1.57	51.53		SBERT-XG	3.56	3.21	33.20
	TF-IDF-SE	2.53	2.00	30.70		TF-IDF-SE	3.58	3.47	28.41
DC	SPERT	1.15	0.64	75.08	XD	SPERT	1.38	1.02	71.61
	SBERT-XG	2.15	1.90	19.61		SBERT-XG	1.91	1.57	67.23
	TF-IDF-SE	0.68	1.00	39.94		TF-IDF-SE	2.01	2.00	20.82
BB	SPERT	0.46	0.41	71.25	TD	SPERT	2.41	2.18	69.23
	SBERT-XG	1.23	1.45	56.32		SBERT-XG	3.63	2.31	48.83
	TF-IDF-SE	1.02	1.13	63.84		TF-IDF-SE	5.05	5.00	3.95
CV	SPERT	1.68	1.23	81.24	TE	SPERT	0.36	0.20	88.71
	SBERT-XG	2.51	1.69	75.21		SBERT-XG	1.91	1.72	63.31
	TF-IDF-SE	4.04	1.00	44.15		TF-IDF-SE	0.97	1.00	33.95

The use of transformer-based embeddings allowed SPERT to capture deep semantic relationships within the user stories, leading to more accurate predictions. This improvement is visualized in Fig. 7, which illustrates the lower error rates of SPERT compared to traditional NLP-based models across multiple projects.

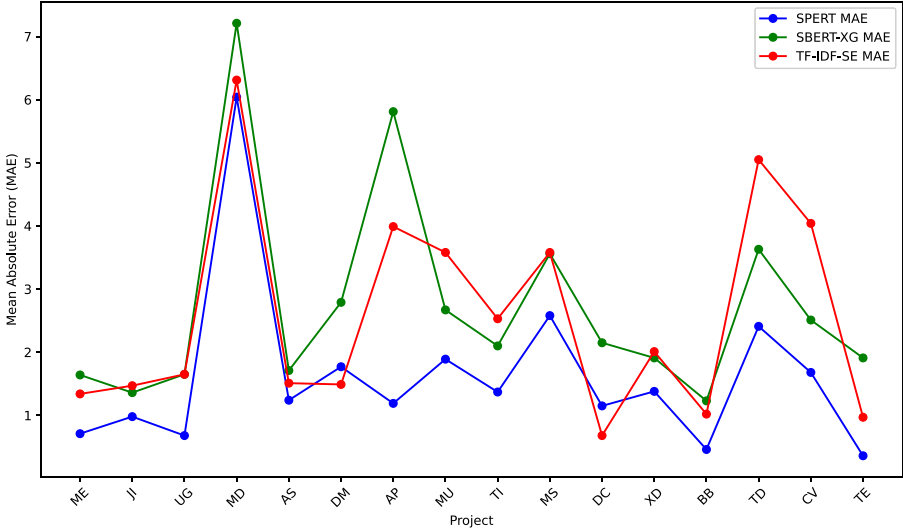


Fig. 7. Performance comparison (MAE) between SPERT and traditional NLP techniques.

RQ4: SPERT in Cross project estimation

The cross-project estimation results are presented in Table 5. SPERT demonstrated superior generalization capabilities compared to the baseline models in all cross-project scenarios. For example, when trained on the Mesos project and tested on the Usergrid project, SPERT achieved an MAE of 0.84, while SBERT-XG recorded an MAE of 1.31.

Figure 8 illustrates the cross-project performance, highlighting that SPERT consistently achieved lower error rates compared to the baseline models. The A12 effect size values, which were consistently above 0.7, further validate the robustness of SPERT's generalization capabilities, making it suitable for Agile environments with varying project characteristics.

8.1. Comparative analysis and limitations of SPERT

While SPERT demonstrates superior performance over baseline models, it is essential to acknowledge its limitations and provide a balanced view of its capabilities.

Training Complexity: SPERT's integration of transformer-based embeddings and RL significantly increases the complexity of its training process compared to traditional methods. The multi-stage training pipeline — involving fine-tuning

Table 5. MAE on cross-project estimation and comparison of SPERT and baselines using Wilcoxon test and A12 effect size (in brackets).

Source	Target	SPERT	Baseline	Wilcoxon (p -value)	Effect (A12)
(i) Within-repository					
ME	UG	0.71	1.23	0.001	[0.78]
UG	ME	0.68	1.22	0.012	[0.72]
AS	AP	1.24	1.90	0.001	[0.74]
AS	TI	1.37	2.56	0.001	[0.76]
AP	AS	1.19	2.35	0.051	[0.66]
AP	TI	1.52	2.80	0.003	[0.60]
MU	MS	1.89	3.02	0.041	[0.65]
MS	MU	2.14	2.89	0.030	[0.62]
Avg		1.34	2.12		[0.69]
(ii) Cross-repository					
AS	UG	1.57	2.14	0.004	[0.61]
AS	ME	1.78	2.50	0.022	[0.55]
MD	AP	4.57	6.10	0.001	[0.60]
MD	TI	5.45	7.14	0.097	[0.54]
MD	AS	5.70	6.88	0.001	[0.62]
DM	TI	2.25	3.89	0.001	[0.64]
UG	MS	4.00	4.45	0.005	[0.55]
ME	MU	2.68	3.10	0.015	[0.60]
Avg		3.00	4.14		[0.59]

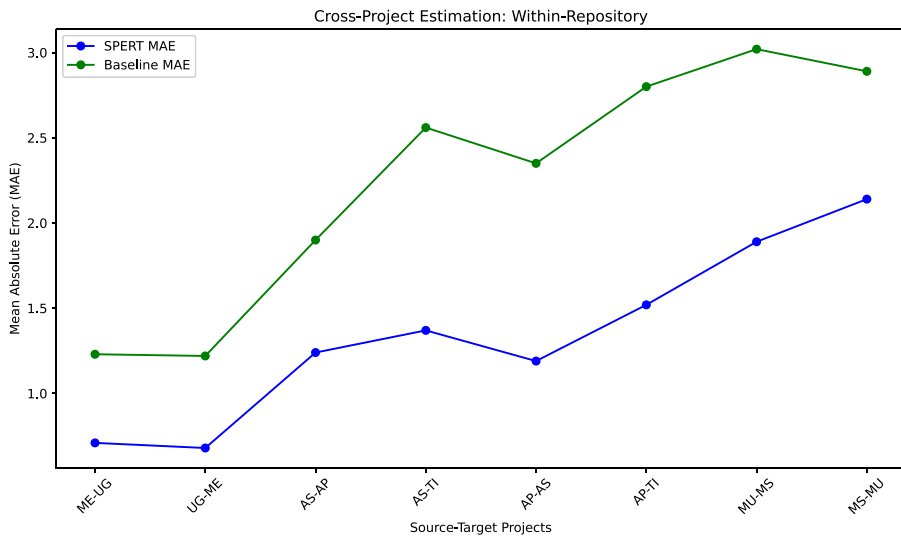


Fig. 8. Cross-project estimation performance of SPERT compared to baseline models.

BERT embeddings and optimizing the PPO algorithm — requires substantial computational resources and time. For instance, training SPERT on large datasets like the ones used in this study can take several hours on high-performance GPUs, making it less feasible for teams with limited computational capacity.

Computational Requirements: The use of BERT embeddings and RL introduces a high computational overhead. During training, SPERT requires multiple passes through the transformer encoder and iterative updates to the policy network. Additionally, memory usage is significantly higher due to the storage of embedding representations and RL buffers, which may limit its applicability in environments with constrained hardware.

Deployment Challenges: Deploying SPERT in real-world Agile environments may pose challenges due to its complexity. The model's reliance on RL makes it sensitive to the quality and consistency of feedback. For example, noisy or biased project data can adversely affect the model's performance, requiring careful pre-processing and monitoring during deployment.

Comparison with Existing Models: Compared to simpler baseline models like SBERT-XG or TF-IDF-SE, SPERT offers superior accuracy but at the cost of increased training time and resource usage. For example, while SBERT-XG can provide reasonably accurate predictions with minimal computational overhead, SPERT's more sophisticated architecture ensures better generalization across diverse projects but demands greater effort in training and tuning.

9. Conclusion


In this paper, we introduced SPERT, a novel model that leverages transformer-based embeddings and RL for more accurate and adaptive story point estimation in Agile software development. Our evaluation demonstrated that SPERT outperforms traditional ML methods and state-of-the-art deep learning models across multiple metrics, including MAE, MdAE and SA. The results highlight the strength of combining semantic text representations through transformers with adaptive prediction mechanisms via RL, ensuring robust performance even in cross-project estimation scenarios.

For future work, we plan to explore further enhancements to SPERT by incorporating additional contextual features, such as team performance metrics or sprint velocities, to improve predictive accuracy. Additionally, experimenting with more advanced RL algorithms and integrating external project management data could offer deeper insights into how evolving project dynamics affect story point estimation. Last, expanding the model's applicability to different domains beyond software development may help generalize SPERT's utility in broader project management contexts.


Acknowledgment


The authors extend their appreciation to King Saud University, Riyadh, Saudi Arabia, for funding this work through Researchers Supporting Project number (RSPD2024R704).


ORCID


Waleed Younas  <https://orcid.org/0000-0002-7858-1779>

Rui Chen  <https://orcid.org/0009-0007-2324-0030>

Jing Zhao  <https://orcid.org/0000-0001-8529-3399>

Tahreem Iqbal  <https://orcid.org/0009-0003-6184-5736>

Mohamed Sharaf  <https://orcid.org/0000-0001-6722-8366>

Azhar Imran  <https://orcid.org/0000-0003-3598-2780>

References

1. K. Schwaber, *Agile Project Management with Scrum* (Microsoft Press, 2004).
2. S. Group, Chaos Report 2015, Standish Group International (2015).
3. K. Beck, M. Beedle, A. Van Bennekum, A. Cockburn, W. Cunningham, M. Fowler, J. Grenning, J. Highsmith, A. Hunt, R. Jeffries, J. Kern, B. Marick, R. C. Martin, S. Mellor, K. Schwaber, J. Sutherland and D. Thomas, Manifesto for agile software development (2001).
4. M. Bloch, S. Blumberg and J. Laartz, Delivering large-scale it projects on time, on budget, and on value, McKinsey & Company and University of Oxford Joint Report (2012).
5. B. Boehm, *Software Engineering Economics* (Prentice Hall, Upper Saddle River, NJ, USA, 1981).
6. A. J. Albrecht, Measuring application development productivity, in *Proc. Joint SHARE/GUIDE/IBM Application Development Symposium*, 1979, pp. 83–92.
7. N. Haugen, Planning poker or how to avoid being victimized by random numbers, XPlanner Blog (2006).
8. K. S. Rubin, *Essential Scrum: A Practical Guide to the Most Popular Agile Process* (Addison-Wesley, 2012).
9. M. Abadeer and M. Sabetzadeh, Machine learning-based estimation of story points in agile development: Industrial experience and lessons learned, in *Proc. 2021 IEEE 29th Int. Requirements Engineering Conf. Workshops*, 2021, pp. 106–115.
10. A. Anand, J. Kaur, O. Singh and O. Alhazmi, Optimal sprint length determination for agile-based software development, *Comput. Mater. Continua* **68** (2021) 3693–3712.
11. B. Yalciner, K. Dincer, A. G. Karacor and M. O. Efe, Enhancing agile story point estimation: Integrating deep learning, machine learning, and natural language processing with sbert and gradient boosted trees, *Appl. Sci.* **14**(16) (2024) 7305.
12. N. Reimers and I. Gurevych, Sentence-bert: Sentence embeddings using siamese bert-networks, in *Proc. 2019 Conf. Empirical Methods in Natural Language Processing*, 2019, pp. 3982–3992.
13. C. Bentéjac, A. Csörgő and G. Martínez-Muñoz, A comparative analysis of gradient boosting algorithms, *Artif. Intell. Rev.* **54** (2021) 1937–1967.
14. W. Xu, S. Wang and J. Chen, Lhc-se: Learning to estimate software effort using deep learning models, *J. Softw.: Evol. Process* **32** (2020) e2290.
15. Y. Gao, X. Zhang and L. Yu, Deep learning-based software effort estimation: A survey, *J. Syst. Softw.* **170** (2020) 110717.
16. M. Choetkiertikul, H. K. Dam, T. Tran and A. Ghose, A deep learning model for estimating story points, *IEEE Trans. Softw. Eng.* **45**(7) (2018) 637–656.
17. Z. Li, B. Wu and X. Wang, A systematic literature review of term weighting schemes for text classification, *Expert Syst. Appl.* **106** (2018) 1–11.

18. J. Devlin, M.-W. Chang, K. Lee and K. Toutanova, Bert: Pre-training of deep bidirectional transformers for language understanding, in *Proc. 2019 Conf. North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, 2019, pp. 4171–4186.
19. Y. Liu, M. Ott, N. Goyal, J. Du, M. Joshi, D. Chen, O. Levy, M. Lewis, L. Zettlemoyer and V. Stoyanov, Roberta: A robustly optimized bert pretraining approach, arXiv:1907.11692 (2019).
20. A. Sousa, D. Veloso, H. Gonçalves, J. Faria, J. Mendes-Moreira, R. Graça, D. Gomes, R. Castro and P. Henriques, Applying machine learning to estimate the effort and duration of individual tasks in software projects, *IEEE Access* **11** (2023) 89933–89946.
21. R. S. Sutton and A. G. Barto, *Reinforcement Learning: An Introduction* (MIT Press, 2018).
22. J. Schulman, F. Wolski, P. Dhariwal, A. Radford and O. Klimov, Proximal policy optimization algorithms, arXiv:1707.06347 (2017).
23. M. Adnan, M. Afzal and K. Asif, Ontology-oriented software effort estimation system for e-commerce applications based on extreme programming and scrum methodologies, *Comput. J.* **62** (2019) 1605–1624.
24. Q. Jiang, Z. Wang, L. Zhang and X. Xia, Automated story point estimation using deep learning techniques, *J. Syst. Softw.* **174** (2021) 110901.
25. D. M. Fernandez, A. Lochmann and S. Wilder, Modeling requirements for agile software projects, *Requir. Eng.* **24** (2019) 175–192.
26. W. Xu, Z. Li, S. Wang and H. Wang, Predicting story points from text using cnn-lstm networks, in *Proc. 26th Int. Requirements Engineering Conf.*, 2018, pp. 11–20.
27. Y. Matsumoto, S. Morikawa, H. Hata and K. Matsumoto, Story point estimation based on developers' mental effort and its influencing factors, in *2020 IEEE/ACM Int. Conf. Software and System Process*, 2020, pp. 11–20.
28. J. Zhang, F. Huang and C. Ding, Learning to estimate story points: A comparison of classical and deep learning methods, *Inf. Softw. Technol.* **107** (2019) 78–89.
29. S. Wang, L. Li and H. Zhang, Software effort estimation based on deep learning, in *Proc. 2020 Int. Conf. Computer Science, Communication and Information Systems*, 2020, pp. 1–5.
30. S. Rathore and D. Kumar, Story point estimation using machine learning: A case study, *Int. J. Comput. Appl.* **178** (2019) 10–14.
31. D. McQuillin, O. Carter and K. Singh, Effort estimation techniques in software development: A comparative analysis, in *2020 IEEE Int. Conf. Software Engineering*, 2020, pp. 1052–1061.
32. H. Tawosi, D. Di Stefano and M. Di Penta, Predicting story points using developer activity and linguistic features, *Inf. Softw. Technol.* **126** (2020) 106348.
33. Y. Pan, W. Xu and J. Chen, A transformer-based framework for predicting software development effort, *J. Softw.: Evol. Process* **33** (2021) e2345.
34. D. M. Fernandez, A. Lochmann and S. Wilder, Improving story point estimation with machine learning: An empirical study, in *Proc. 2021 Int. Conf. Software Engineering*, 2021, pp. 234–244.
35. H. M. Faris, H. Faris, S. Ibrahim, M. Aloqaily and M. Otair, Text analytics and story points estimation using nlp techniques, *IEEE Access* **9** (2021) 73170–73185.
36. Z. Mao, J. Zhao, Y. Wang and W. Jin, Machine learning-based software effort estimation: A comparative study, *Appl. Soft Comput.* **83** (2019) 105657.
37. J. Gao, X. Zhao and J. Wang, Cross-project software effort estimation using transfer learning and deep learning, *J. Syst. Softw.* **176** (2021) 110946.

38. S. Levine, Reinforcement learning and control as probabilistic inference: Tutorial and review, arXiv:1805.00909 (2018).
39. X. Xia, Q. Wang, W. Zhou and X. Gao, Improving cross-project effort estimation with fine-tuned pretrained language models, *Inf. Softw. Technol.* **132** (2021) 106495.
40. Y. Pan, H. Wang and H. Huang, Contextual story point estimation using pretrained language models, *Inf. Softw. Technol.* **130** (2021) 106482.
41. C. Niu, Y. He and L. Zhao, Adaptive software development effort estimation based on deep neural networks, *Expert Syst. Appl.* **174** (2021) 114767.
42. M. Schalken, M. van Kuilenburg, J. J. van Wijk and A. C. Telea, Effort estimation for software projects: A comparison of machine learning techniques, in *2023 IEEE/ACM 45th Int. Conf. Software Engineering*, 2023, pp. 526–536.
43. V. Patel, G. Singh and R. Sharma, Sentence-bert for story point estimation: A deep learning approach, *Inf. Softw. Technol.* **124** (2020) 106357.
44. S. Pargaonkar, A comprehensive research analysis of software development life cycle (sdlc) agile & waterfall model advantages, disadvantages, and application suitability in software quality engineering, *Int. J. Sci. Res. Publ.* **13** (2023) 345–358.
45. N. Tran, T. Tran and N. Nguyen, Leveraging AI for enhanced software effort estimation: A comprehensive study and framework proposal, arXiv:2402.05484 (2024).
46. N. Srivastava, G. Hinton, A. Krizhevsky, I. Sutskever and R. Salakhutdinov, Dropout: A simple way to prevent neural networks from overfitting, *J. Mach. Learn. Res.* **15**(1) (2014) 1929–1958.
47. I. Loshchilov and F. Hutter, Decoupled weight decay regularization, arXiv:1711.05101 (2017).
48. Z. Chen, Y. Li, B. R. Wang and S. Wang, Sequencer: Sequence-to-sequence learning for end-to-end program repair, *IEEE Trans. Softw. Eng.* **47**(9) (2019) 1943–1959.
49. F. Sarro, A. Petrozziello and M. Harman, Multi-objective software effort estimation, *Autom. Softw. Eng.* **23**(3) (2016) 462–494.
50. M. Shepperd, S. G. MacDonell, G. Kadoda and B. Kitchenham, Evaluating prediction systems in software project estimation, *Inf. Softw. Technol.* **54**(8) (2012) 820–827.
51. V. Tawosi, F. Sarro, A. Petrozziello and M. Harman, Multi-objective software effort estimation: A replication study, *IEEE Trans. Softw. Eng.* **48** (2021) 3185–3205.