

Bug-Report-Driven Fault Localization: Industrial Benchmarking and Lessons Learned at ABB Robotics

Pernilla Hall
ABB Robotics
Sweden
pernilla.hall@e.abb.com

Riccardo Rubei
Mälardalen University
Sweden
riccardo.rubei@mdu.se

Anton Ununger
ABB Robotics
Sweden
anton.ununger@u@se.abb.com

Alessio Bucaioni
Mälardalen University
Sweden
alessio.bucaioni@mdu.se

Abstract

Software quality assurance remains a major challenge in industrial environments, where large-scale and long-lived systems inevitably accumulate defects. Identifying the location of a fault is often time-consuming and costly, particularly during maintenance phases when developers must rely primarily on textual bug reports rather than complete runtime or code-level context. In this study, we investigated whether artificial intelligence can support fault localization using only the natural-language content of bug reports. By relying exclusively on textual information, our approach requires no access to source code, execution traces, or static analysis artifacts, making it directly deployable within existing industrial maintenance workflows.

We framed fault localization as a supervised text classification problem and evaluated three traditional machine learning models (Logistic Regression, Support Vector Machine, and Random Forest) and two fine-tuned transformer-based language models (RoBERTa-Base and Distil-RoBERTa). Our evaluation used proprietary data from ABB Robotics in Västerås, Sweden, comprising approximately five years of resolved industrial bug reports, each linked to its verified code fix. This setting allowed us to assess model effectiveness under realistic industrial constraints.

Our results showed that traditional models using term frequency-inverse document frequency features consistently outperformed the fine-tuned language models on this dataset, while data augmentation substantially improved Random Forest performance. These findings challenge the assumption that transformer-based models universally outperform classical approaches in industrial contexts with domain-specific data. Overall, we demonstrated that historical bug reports can be systematically leveraged for text-based, artificial intelligence-assisted fault localization, providing a scalable, low-cost, and empirically grounded complement to traditional debugging practices in industry.

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

EASE 2026, Glasgow, United Kingdom

© 2026 Copyright held by the owner/author(s). Publication rights licensed to ACM.
ACM ISBN 978-x-xxxx-xxxx-x/YYYY/MM
<https://doi.org/10.1145/nnnnnnn.nnnnnnn>

ACM Reference Format:

Pernilla Hall, Anton Ununger, Riccardo Rubei, and Alessio Bucaioni. 2026. Bug-Report-Driven Fault Localization: Industrial Benchmarking and Lessons Learned at ABB Robotics. In *Proceedings of The 30th International Conference on Evaluation and Assessment in Software Engineering (EASE 2026)*. ACM, New York, NY, USA, 9 pages. <https://doi.org/10.1145/nnnnnnn.nnnnnnn>

1 Introduction

Modern software-intensive systems are large, interconnected, and continuously evolving. As a result, defects are inevitable and costly to diagnose and fix [19]. A central bottleneck is fault localization, that is, identifying where a fault originates, which is time-consuming, expertise-intensive, and often performed under incomplete information [24]. In industrial maintenance settings, developers frequently rely primarily on textual bug reports during triage, without guaranteed access to execution traces, dynamic logs, or detailed architectural context. At the same time, modern bug tracking systems accumulate large volumes of historical reports linked to verified fixes. When these reports contain sufficient diagnostic signals, the mapping between report text and fix location provides a practical foundation for Machine Learning (ML)-assisted fault localization.

In this work, we conducted a large-scale industrial study of *text-only fault localization*: given a bug report (title and description), we predicted a ranked list of likely fault locations using only natural-language information. To ensure actionable recommendations, we predicted at the subfolder (component) level rather than at the individual file level. This design mirrors how developers navigate large industrial codebases and mitigates label-space sparsity. Prior research has primarily evaluated fault localization on open-source systems and often assumes access to auxiliary artifacts such as execution traces or call graphs [6, 10, 11, 16, 28]. Such assumptions rarely hold in industrial maintenance contexts. Moreover, evaluations of Large Language Models (LLMs) on open-source corpora risk overlap with pretraining data, potentially inflating performance estimates. To assess feasibility under realistic industrial and confidentiality constraints, we analyzed a proprietary, large-scale ABB Robotics system in Västerås, Sweden, linking five years of resolved bug reports to their verified code changes.

We compared three traditional ML models, Logistic Regression (LR), Support Vector Machine (SVM), and Random Forest (RF),

against two fine-tuned transformer-based language models, RoBERTa-Base and Distil-RoBERTa. The traditional models used term frequency-inverse document frequency (TF-IDF) representations and sentence embeddings, whereas the LLMs operated directly on tokenized reports. We applied systematic text preprocessing and class-imbalance-aware data augmentation. We evaluated all models using ranking-centric metrics (Top- k Accuracy, Recall@ k , MAP, and MRR) that reflect developer decision workflows. Our study addressed the following research questions:

RQ1: How do Random Forest, Logistic Regression, and Support Vector Machine perform for text-only fault localization in a proprietary industrial software system?

RQ2: How does a fine-tuned RoBERTa model (Base/Distil) compare to traditional ML models under the same industrial constraints?

Our results showed that LR and SVM with TF-IDF features consistently outperformed more complex approaches on non-augmented data. Data augmentation substantially improved RF, and although both LLMs benefited from augmentation, they did not consistently surpass the strongest TF-IDF baselines; Distil-RoBERTa generalized more robustly than RoBERTa-Base. These findings provide concrete evidence that, in data-constrained and confidentiality-bound industrial settings, lightweight and interpretable models can match or exceed transformer-based alternatives. Despite limitations, including the single-system scope, variable report quality, and pronounced label imbalance, the evidence demonstrated that text-only ML models can reliably narrow developers' initial search space and support cost-effective triage decisions.

This paper makes the following contributions:

- A large-scale industrial benchmark of text-only fault localization. We evaluated ML-assisted localization on a proprietary ABB Robotics system under realistic industrial and confidentiality constraints, mapping bug reports to verified component-level fault locations.
- A controlled comparison of lightweight models and fine-tuned LLMs. We systematically benchmarked LR, SVM, and RF against RoBERTa-Base and Distil-RoBERTa using ranking metrics aligned with developer workflows.
- Decision-oriented guidance for industrial adoption. We identified when simple TF-IDF-based models suffice and when the additional complexity of LLMs yields limited marginal benefit, thereby informing evidence-based technology choices.
- Reproducible artifacts for technology transfer. We provided scripts and configuration files as a replication package (excluding proprietary data) to enable application in other industrial contexts.

Section 2 positions our study within related work. Section 3 describes the overall pipeline, and Section 4 details the experimental setup and implementation. Section 5 presents the results, which we discuss in Section 6. Section 7 concludes the paper with final remarks and future research directions.

2 Related Work

Fault localization has been studied through similarity-based, historical co-change, program-structure-aware, and deep learning approaches, including statistic-based methods [17, 18, 23], program

state-based approaches [7, 30, 31], ML-based techniques [4, 5, 25], and model-based solutions [1, 8, 9]. Most evaluations relied on open-source datasets and assumed auxiliary artifacts such as execution traces, call graphs, stack traces, or repository navigation, which were often unavailable in industrial maintenance. In addition, LLM evaluations on public corpora risked overlap with pretraining data and could inflate performance estimates. In contrast, we evaluated text-only fault localization under realistic industrial constraints using proprietary ABB Robotics data spanning five years of resolved bug reports linked to verified fixes.

Several approaches exploited historical and structural relationships, often through graph-based representations. Chen et al. [6] modeled co-fixed files as an implicit social network and applied PageRank, reporting Top-10 accuracies of 32–41%. Li et al. [16] proposed LaProb, which connected reports and files via textual similarity and call graphs and propagated labels from known fixes, achieving Top-1 accuracy between 14–68% across projects. Gharibi et al. [11] combined textual and semantic similarity, stack traces, keyword matching, and a historical multi-label Naive Bayes classifier, achieving average Top-1 \approx 52%, Top-10 \approx 78%, and MAP between 0.32 and 0.65 on three open-source systems. These results highlighted the value of combining multiple artifacts, but they depended on inputs that were not consistently accessible in industry.

Recent work explored LLM-based fault localization [20, 22, 26, 32]. Kang et al. [14] proposed AUTOFL, where an LLM navigated a repository through function calls to retrieve code context, reporting Top-1 between \approx 29% and 53%. Such approaches required interactive repository access and often depended on cloud APIs, which raised confidentiality, latency, and deployment concerns. FlexFL [27] introduced a two-stage framework for identifying root causes in source code and supported lightweight models within a multi-agent architecture. ChatDL [29] addressed change-based defect localization in flexible manufacturing software and showed that combining LLMs with information retrieval outperformed standalone GPT-based solutions. In contrast, we evaluated locally executable, text-only models without traces, call graphs, or repository interaction at inference time.

Traditional ML models have also performed well on bug report classification. Andrade et al. [3] benchmarked SVM, LR, RF, k-NN, and Naive Bayes with TF-IDF and found SVM, LR, and RF generally strongest while remaining sensitive to class imbalance. Köksal and Tekinerdogan [15] demonstrated industrial feasibility in a bilingual setting, achieving a weighted F1 of 72% with a linear SVM. Other work categorized reports into technical and functional classes (e.g., Concurrency, Memory) and showed that ensembles could outperform single models despite noisy reports and class difficulty [12]. Transformer-based classifiers such as CatIss [13] further improved contextual categorization when sufficient data and metadata were available.

Unlike classification-focused studies, we addressed a more demanding *multi-label ranking* problem by predicting likely components (subfolders) within a large and highly imbalanced label space. We emphasized actionable precision at top- k , and we compared lightweight TF-IDF baselines with fine-tuned LLMs under identical industrial constraints to provide guidance for deploying ML-assisted fault localization in real maintenance workflows.

3 Experimental Pipeline

Our experimental pipeline comprised six stages: data collection, preprocessing, feature extraction, model selection, training, and evaluation.

3.1 Data Collection

We analyzed approximately five years of resolved bug reports from a large-scale ABB Robotics system (Västerås, Sweden). Each report was systematically linked to the source-code commits that fixed the issue, enabling supervised learning that maps natural-language descriptions to concrete maintenance locations in a complex, modular, and long-evolving industrial product. Using proprietary data reduced the risk of LLM pretraining overlap and provided a realistic estimate of industrial generalization. Due to confidentiality constraints, raw reports and code changes cannot be released; however, we document the full extraction and evaluation procedure to support replication on comparable industrial datasets.

Input Features. For each report, we extracted two textual inputs: (i) the bug title, which typically contains domain-specific keywords indicating the affected subsystem, and (ii) the bug description, which provides contextual information such as reproduction steps, observed behavior, and occasional log fragments. We deliberately excluded non-textual metadata (e.g., severity, status, assignee) to preserve a realistic text-only deployment scenario and prevent models from exploiting process artifacts unrelated to localization.

Target Labels. Target labels represented component-level fault locations at subfolder granularity. For each resolved report, we extracted modified file paths from linked commits, normalized them relative to the repository root, and mapped each file to its parent subfolder. The resulting set of subfolders formed a potentially multi-label ground truth. We chose this abstraction for three reasons: (i) architectural coherence, as subfolders align with modules or services; (ii) developer actionability, since teams navigate code at this granularity; and (iii) learning tractability, because aggregating to subfolders substantially reduces label-space sparsity compared to file-level prediction in large industrial systems.

3.2 Data Preprocessing

We applied a deterministic preprocessing pipeline to standardize textual inputs. We removed reports lacking descriptions or linked fixes to ensure reliable supervision. We then applied lowercasing, punctuation removal, stop-word removal, de-camelcasing of identifiers (e.g., `RobotArmController` → “robot arm controller”), lemmatization, and tokenization. These steps reduced stylistic variability while preserving semantically relevant content for both traditional ML and transformer-based models.

Feature Engineering. For traditional ML models, we used two representations: (i) TF-IDF Bag-of-Words vectors to capture lexical salience, and (ii) sentence embeddings generated with the sentence-transformers library to encode contextual similarity beyond surface forms. This enabled a controlled comparison between lexical and semantically enriched feature spaces. For transformer-based models, we tokenized reports using each model’s native tokenizer. We

encoded targets as multi-hot vectors over subfolders, supporting multi-label prediction when a bug affected multiple components.

3.3 Algorithm Selection

Based on prior work in ML-based bug report analysis and fault localization [2, 3, 12, 15, 33], and considering on-premise hardware constraints, we compared three traditional models and two transformer-based models:

- RF: robust to overfitting and effective in high-dimensional spaces.
- LR: linear, interpretable, and a strong baseline for sparse text features.
- SVM: well-suited for high-dimensional and sparse text representations.
- RoBERTa-Base and Distil-RoBERTa: encoder-only pretrained models evaluated to balance predictive performance and computational cost.

We wrapped traditional models with `OneVsRestClassifier` to handle multi-label prediction and produce ranked outputs suitable for developer triage.

3.4 Model Training

We partitioned the dataset into training, validation, and test splits (70/20/10). We used the training set for model fitting, the validation set for hyperparameter tuning and early stopping (particularly during LLM fine-tuning), and the test set for a single, unbiased evaluation of generalization. We randomized splits at the bug-report level and applied multi-label stratification to preserve label distributions across partitions, preventing leakage and ensuring robust performance estimates.

3.5 Evaluation and Comparative Analysis

Our objective was to reduce developers’ search space by returning a ranked list of likely subfolders. We therefore evaluated models on the held-out test set using ranking-oriented, threshold-independent metrics:

- Top- k Accuracy ($\text{Hit}@k$) measures the proportion of reports for which at least one true label appears among the top- k predictions ($k \in \{1, 3, 5, 10\}$).
- Recall@ k measures the fraction of true labels retrieved within the top- k predictions, averaged across reports.
- Mean Average Precision (MAP) evaluates ranking quality by averaging precision at each rank where a relevant label appears, rewarding early correct predictions.
- Mean Reciprocal Rank (MRR) captures how early the first relevant subfolder appears in the ranked list.

These metrics emphasize early precision and overall retrieval effectiveness in a multi-label setting and are standard in fault localization and bug-report analysis [11, 16]. By evaluating all models under identical industrial constraints, we ensured a fair and practically meaningful comparison.

Table 1: Synthesized example of the CSV file with extracted bugs.

Date	Bug ID	Bug Title	Prio.	Description	Paths
2021-04-02 13:57	123456	Project Validation Error	2	Warning message when using non UTF-8 character in project name...	../././././File-X.cs
2022-01-22 14:21	293874	Fatal Error When Saving	1	Fatal error when saving a project with non-unique name...	../././././File-Y.cs, ../././././File-Z.cs

4 Experimental Setup

This section describes the concrete implementation of our experimental pipeline, including environment configuration, data extraction and preparation, feature engineering, model training and tuning, and final evaluation. We ran all experiments on a Windows 11 machine using Python 3.10.9 and Visual Studio Code. We used pandas and numpy for data processing, scikit-learn for training and evaluating the traditional models, and transformers for fine-tuning and inference with RoBERTa-based models. To support reproducibility, we provide a complete, version-pinned dependency list in the replication package in Section 8 .

4.1 Data Collection

As outlined in Section 3, the dataset originated from a proprietary ABB Robotics software project (Västerås, Sweden). We implemented a six-step acquisition pipeline: (1) retrieve closed bug reports, (2) retain only reports linked to commits, (3) extract and aggregate modified file paths, (4) map files to subfolder-level components, (5) filter large and diffuse fixes, and (6) remove rare labels to mitigate extreme sparsity.

First, we queried the work-item database using WIQL to retrieve items of type Bug in state Closed within the project scope. This yielded 1,867 reports spanning approximately five years. We then queried work-item relations and retained only bugs linked to at least one Commit artifact. For the remaining reports, we extracted all modified file paths from linked commits and aggregated them per bug, producing a structured CSV file with 802 unique bug reports (see Table 1).

Next, we mapped each modified file path to its parent subfolder to define component-level target labels. We enumerated all unique file paths and their frequencies across commits, and we manually curated a mapping in Excel by adding a Label column corresponding to the architecturally meaningful subfolder path (see the synthesized example in Table 2). This step ensured alignment with domain-specific module boundaries rather than relying on purely syntactic path prefixes.

Table 2: Synthesized example of the CSV file with mapping between files and their corresponding subfolder.

File	Occ.	Label
Root/Project/Folder/SubFolder-X/File-A.cs	25	Root/Project/Folder/SubFolder-X/
Root/Project/Folder/SubFolder-Y/File-B.cs	22	Root/Project/Folder/SubFolder-Y/

We applied the curated mapping to append a label_list column to each bug entry, capturing the unique subfolders modified for that report. We then applied two filtering steps to improve signal quality and learnability.

- Large, diffuse fixes. We removed bugs linked to more than five distinct subfolders. Exploratory inspection showed that

these cases typically corresponded to cross-cutting refactorings or broad maintenance changes, which dilute localization signals and distort ranking-based evaluation.


- Rare labels. We removed subfolder labels with fewer than 10 occurrences across the corpus and discarded bugs whose label lists became empty. This step reduced extreme class imbalance and ensured that each retained label had sufficient training support.

After filtering, the final dataset comprised 660 unique bug reports and 31 subfolder labels. The mean number of labels per bug was 1.56. The distribution of labels per report is shown in Table 3.

Table 3: Distribution of labels per bug report.

Number of Labels	1	2	3	4	5
Number of Bugs	428	138	61	20	13

Prior empirical studies have shown that defect distributions often follow the Pareto principle, where a small fraction of components accounts for most faults [21]. Our dataset exhibited this pattern. As illustrated in Figure 1, the label distribution was strongly skewed: the five most frequent subfolders occurred substantially more often than the remaining components, producing a pronounced long tail and significant class imbalance that motivated our imbalance-aware modeling strategy.

**Figure 1: Visualization of label imbalance in the dataset.**

4.2 Text Preprocessing

We applied a deterministic natural-language preprocessing pipeline before splitting the dataset and training models. We concatenated the Bug Title and Description fields and normalized the resulting text using the following operations:

- remove Azure DevOps template phrases (e.g., “Detailed Description:”) using regular expressions,

Table 4: Optimal Random Forest hyperparameters for all datasets and features.

Parameter	Values tested	Orig (Emb.)	Orig (TF)	RS Full	RS Targ.	SR Full	SR Targ.
TF-IDF max features	1000, 5000, 10000, None	N/A	1000	1000	1000	5000	1000
N-gram range	(1,1), (1,2), (1,3)	N/A	(1,1)	(1,1)	(1,1)	(1,1)	(1,2)
Min document frequency	1, 2, 3	N/A	2	2	2	2	2
Number of trees	5, 10, 20, 100	20	20	100	100	100	100
Tree max depth	None, 20, 100	None	None	None	None	None	None
Min samples split	2, 5, 10	2	2	5	5	2	2
Min samples leaf	2, 5, 10, 20	20	10	2	2	2	2
Class weight	balanced subsample	balanced subsample	balanced subsample	balanced subsample	balanced subsample	balanced subsample	balanced subsample

- de-camelcase identifiers using regular expressions,
- case fold to lowercase,
- strip punctuation using regular expressions,
- tokenize and remove stop words using NLTK’s standard English list,
- lemmatize tokens using NLTK’s WordNetLemmatizer.

The processed dataset contained the columns *processed_text* and *label_list*, and we used it as input to dataset splitting and modeling.

4.3 Dataset Splitting and Label Binarization

To ensure reproducibility and a like-for-like comparison across models, we performed a single randomized split (seed = 42) of the preprocessed dataset into 70% training, 20% validation, and 10% test. Given the multi-label setting and class imbalance, we used scikit-multilearn’s `iterative_train_test_split` to apply iterative stratification and approximately preserve per-label frequencies across partitions, which ensured that all labels remained represented despite the limited dataset size. After splitting, we binarized target labels (subfolder lists) into multi-hot vectors using scikit-learn’s `MultiLabelBinarizer`. We fitted the binarizer on the training labels only to prevent leakage, and then applied it to the validation and test splits.

4.4 Data Augmentation

To mitigate limited data and pronounced class imbalance, we applied text-based augmentation only to the training split to increase diversity and sample counts without contaminating validation or test data. We used two techniques: (i) WordNet-based synonym replacement and (ii) random word swap, both producing semantically similar but syntactically varied reports to improve generalization.

We evaluated two augmentation strategies: (i) full augmentation of the entire training set and (ii) targeted augmentation for under-represented labels (occurrence < 25). We used an augmentation factor of 1, generating one augmented sample per original and doubling the affected subset. This procedure yielded four additional training variants alongside the original split (Section 4.3): full synonym replacement, full random swap, targeted synonym replacement, and targeted random swap. For each variant, we regenerated binarized labels as NumPy arrays using a `MultiLabelBinarizer` fitted on the corresponding training labels, resulting in five training sets in total for model training and benchmarking.

4.5 Feature Extraction

To prepare text into a suitable format of numerical features for traditional ML models, we test and employ two strategies: TF-IDF and sentence embeddings.

4.6 Feature Extraction

We transformed preprocessed text into numerical representations. For traditional ML models, we used TF-IDF vectors and sentence embeddings. For transformer-based models, we applied RoBERTa-style tokenization consistently during fine-tuning and inference.

TF-IDF. We represented each report with scikit-learn’s `TfidfVectorizer`. We tuned `max_features`, `ngram_range`, and `min_df` per model (Section 4.7). To prevent leakage, we fitted the vectorizer on training text only, saved it with `joblib`, and transformed the training, validation, and test splits into sparse feature matrices.

Sentence Embeddings. We encoded reports with `all-mpnet-base-v2` using `sentence-transformers`. Because LR and SVM are sensitive to feature scale, we standardized embeddings with scikit-learn’s `StandardScaler` fitted on training embeddings only, saved the scaler with `joblib`, and applied it to validation and test. The scaled vectors served as inputs to the traditional ML models.

Tokenization. We tokenized reports with the RoBERTa-compatible `AutoTokenizer` (Hugging Face) to produce input IDs and attention masks. We truncated and padded sequences to 256 tokens, which covered most reports in exploratory analysis while keeping compute and memory cost manageable.

4.7 Model Training and Tuning

Traditional ML Pipeline. We evaluated RF, LR, and SVM under a unified pipeline with two text representations: TF-IDF vectors and sentence embeddings. For TF-IDF, we trained a scikit-learn Pipeline (`TfidfVectorizer` + classifier wrapped in `OneVsRestClassifier`) and tuned hyperparameters via grid search with 3-fold iterative stratified cross-validation on the *training* split only. We selected configurations by maximizing cross-validated Mean Average Precision (MAP), refit the best model on the full training set, and saved the trained model and fitted vectorizer with `joblib`. For embeddings, we encoded reports with `all-mpnet-base-v2` (`sentence-transformers`), standardized vectors with a `StandardScaler` fitted on training embeddings only, and applied the same 3-fold iterative stratified CV and MAP-based selection. We refit the selected model on the full training set and saved the trained model and scaler. To analyze augmentation effects, we retrained models on each augmented training variant from Section 4.4 while

Table 5: Optimal hyperparameters for RoBERTa-Base and Distill RoBERTa for all datasets.

Parameter	Values tested	Original		RS Full		RS Targeted		SR Full		SR Targeted	
		Common	Base	Distill R.	Base	Distill R.	Base	Distill R.	Base	Distill R.	Base
Learning rate	1e-6 to 5e-4	2.48e-05	1.13e-4	6.95e-05	1.59e-4	5.21e-05	1.73e-4	3.90e-05	6.92e-05	4.34e-05	1.02e-05
Batch size	4, 8, 16	8	8	16	8	8	16	8	16	16	8
Epochs	20	20	20	20	20	20	20	20	20	20	20
Weight decay	0.0 to 0.3	0.22	0.21	0.02	0.20	0.24	0.27	0.08	0.15	0.09	0.09
Focal Loss Alpha	0.1 to 0.5	0.38	0.50	0.41	0.29	0.30	0.30	0.39	0.17	0.36	0.24
Focal Loss Gamma	1 to 5	4	2	1	1	3	2	1	3	2	5
Max sequence length	256	256	256	256	256	256	256	256	256	256	256
Gradient accumulation steps	2	2	2	2	2	2	2	2	2	2	2

keeping validation and test splits unchanged. Because preliminary experiments showed TF-IDF outperforming embeddings, we focused augmentation analysis on the TF-IDF pipeline for controlled comparisons. Figure 2 summarizes the TF-IDF versus embedding results.

Hyperparameter Selection. We defined model-specific grids and selected hyperparameters based on cross-validated MAP. Table 4 reports the explored configurations and best settings for Random Forest; we provide the full grids and results for all models in the replication package. In Table 4, Original denotes non-augmented training data, SR Full full synonym replacement, RS Full full random swap, SR Targeted synonym replacement for under-represented labels only, and RS Targeted random swap for under-represented labels only.

4.8 Deep Learning Pipeline

We implemented RoBERTa-Base and Distil-RoBERTa using HuggingFace Transformers and Datasets, and we optimized hyperparameters with Optuna. For confidentiality, we pre-downloaded all models and tokenizers and executed the full pipeline locally. The pipeline comprised the following steps:

- load the train/validation/test splits with multi-hot labels and wrap them as Dataset objects,
- tokenize reports with the corresponding RoBERTa tokenizer (truncation and padding to 256 tokens),
- fine-tune with a custom Trainer using Focal Loss to mitigate label imbalance,
- run 20 Optuna trials per model to sample training configurations (e.g., learning rate, batch size, weight decay, and focal-loss parameters), optimizing MAP on the validation split for comparability with traditional models,
- train the final model on the *training* split using the best trial’s configuration, keeping the validation split fixed for selection only, and persist the model, tokenizer, and configuration for evaluation on the held-out test set.

To assess augmentation effects, we repeated the same procedure on each augmented *training* variant (Section 4.4) while keeping validation and test splits unchanged. Hyperparameter ranges and selected optima are reported in Table 5.

4.9 Final Benchmark and Evaluation Scripts

After selecting hyperparameters on the validation split, we ran the final benchmark on the held-out test split. The evaluation script:

- loaded the test set and its binarized labels,
- restored the best saved artifacts per model (traditional models: estimator with TF-IDF vectorizer and, where applicable, StandardScaler; RoBERTa models: fine-tuned model and tokenizer),
- generated per-label scores (probabilities/logits) and ranked subfolders accordingly for each report, and
- computed ranking metrics: Top- k Accuracy ($k \in \{1, 3, 5, 10\}$), Recall@ k , Mean Reciprocal Rank (MRR), and Mean Average Precision (MAP).

We did not retrain or refit any model during this stage. All evaluations used the fixed test split under identical preprocessing. We report the results in Section 5 to support a fair comparison between traditional and deep learning approaches.

5 Results

We present results for the two research questions. We evaluated all configurations once on the held-out test set using the ranking metrics from Section 3.


5.1 Comparison of Traditional ML Models (RQ1)

We benchmarked LR, SVM, and RF in the text-only industrial setting using TF-IDF features. For each model, we report the strongest training regime identified during tuning (Section 4.7) and evaluated on the fixed test split: LR with Full Random Swap (RS), SVM on Original data, and RF with Full Synonym Replacement (SR). This comparison highlights how augmentation interacted with each algorithm under an identical evaluation protocol. Overall (Table 6), LR+TF-IDF (RS Full) delivered the best top-ranked behavior, achieving the highest MRR and tying for Top-1 Accuracy. RF+TF-IDF (SR Full) achieved the highest MAP and Recall@1, indicating the strongest overall ranking quality with augmentation. SVM+TF-IDF (Original) led on Top-5 Accuracy and Recall@5, providing a strong baseline without augmentation. These results showed that TF-IDF remained highly competitive on proprietary industrial data, augmentation benefited RF the most, and LR with RS offered the strongest top-ranked recommendations. In practice, teams can prioritize LR for rapid triage, SVM for stable performance without

Table 6: Best traditional model configurations (test set).

Metric	LR + TF-IDF (RS Full)	SVM + TF-IDF (Original)	RF + TF-IDF (SR Full)
Top-1 Acc.	0.5263	0.5000	0.5263
Top-3 Acc.	0.7632	0.7632	0.7763
Top-5 Acc.	0.8421	0.8553	0.7895
Top-10 Acc.	0.9342	0.9211	0.9079
Recall@1	0.4178	0.4090	0.4265
Recall@3	0.6776	0.6919	0.7018
Recall@5	0.7719	0.7851	0.7390
Recall@10	0.8849	0.8805	0.8717
MAP	0.6109	0.6103	0.6171
MRR	0.6655	0.6526	0.6564

augmentation, or RF when augmentation is feasible and overall ranking quality matters.

**Figure 2: MAP comparison: TF-IDF vs. sentence embeddings for traditional models (test set).**

5.2 Comparison of RoBERTa Models with Traditional Models (RQ2)

We compared the best traditional TF-IDF pipelines with the best fine-tuned RoBERTa variants under identical text-only industrial constraints. As shown in Table 7, neither RoBERTa-Base nor Distil-RoBERTa outperformed the strongest traditional baselines. LR with RS augmentation tied for the best Top-1 Accuracy and achieved the highest MRR, RF with SR augmentation attained the best MAP and Recall@1, and SVM without augmentation led on Top-5 Accuracy and Recall@5. Distil-RoBERTa marginally outperformed RoBERTa-Base on several metrics, but both consistently trailed the TF-IDF baselines on this limited and imbalanced dataset. Overall, the results indicated that lightweight, locally runnable TF-IDF models remained the most effective and deployable option for early-hit component ranking in industrial maintenance.

6 Discussion

In this section, we interpret the results for the two research questions and derive implications for industrial adoption.

LR and SVM with TF-IDF features outperformed RF on the non-augmented dataset, and their near-identical results indicate that linear decision boundaries over sparse lexical features captured most of the useful signal in industrial bug reports. TF-IDF consistently surpassed sentence embeddings, showing that domain-specific terminology and recurring keywords were more predictive

than generalized semantic similarity in this setting. Augmentation produced model-dependent effects. RF benefited substantially from synonym replacement and random swaps, while LR and SVM showed marginal or slightly negative changes, likely because perturbations diluted discriminative terms. These results reinforce prior evidence that linear models with lexical features remain strong baselines for issue-text tasks [3, 15].

Neither RoBERTa-Base nor Distil-RoBERTa surpassed the strongest TF-IDF baselines under identical text-only constraints. Distil-RoBERTa generalized more robustly than RoBERTa-Base on the original data, consistent with expectations for smaller models under limited and imbalanced training conditions. Although augmentation improved both transformers, their absolute performance remained below LR/SVM and competitive RF. Given the dataset size (660 reports, 31 labels) and domain-specific jargon, pretrained language models had limited opportunity to exploit broad semantic knowledge, whereas TF-IDF directly leveraged precise lexical cues aligned with component boundaries. Compared with systems such as AUTOFL [14], which integrate repository navigation and code context, our text-only approach prioritized deployability and confidentiality while maintaining strong ranking quality.

From an industrial perspective, deployability and cost-efficiency are central. TF-IDF pipelines (LR/SVM) and RF with augmentation provided a fully on-prem ranking aid that narrowed the search space without requiring source-code access or external APIs. On the test set, Top-1 around 0.53 (LR/RF) indicated that the first recommendation was correct in over half of cases, while Top-5 up to 0.86 and Recall@5 around 0.79 showed that shortlists typically surfaced at least one relevant component. These levels support use as an assistive triage tool rather than a fully automated locator. A pragmatic adoption path is to start with LR or SVM using TF-IDF and introduce RF with augmentation when label imbalance becomes limiting. Integrating ranked predictions directly into bug trackers or CI dashboards maximizes impact during early investigation. Periodic retraining and monitoring of Top- k and MAP help detect drift and maintain reliability. Overall, carefully tuned lightweight text-based models can deliver substantial practical value under realistic industrial constraints.

6.1 Threats to validity

Our findings were derived from a single proprietary ABB Robotics system, which limits external validity. Results may vary across domains, architectures, programming languages, reporting practices, and alternative label granularities. Replication on additional industrial datasets is therefore necessary to assess generalizability. The dataset size (660 reports) and pronounced class imbalance constrained model capacity, particularly for transformer-based models. These characteristics may have reduced the potential advantage of deep architectures. In addition, confidentiality requirements mandated fully on-premise training and evaluation, which restricted exploration of larger model variants or API-based systems. Construct validity may be influenced by our choice of subfolder-level targets as proxies for actionable components. Other organizations may prefer file-, package-, or service-level labels, which would alter task difficulty, label sparsity, and the interpretation of ranking metrics. Internal validity risks stem from preprocessing, augmentation

Table 7: Best configurations per model family (test set).

Metric	LR+TF-IDF (RS Full)	SVM+TF-IDF (Orig.)	RF+TF-IDF (SR Full)	RoBERTa-Base (RS Targeted)	Distil-RoBERTa (RS Targeted)
Top-1 Acc.	0.5263	0.5000	0.5263	0.4474	0.4605
Top-3 Acc.	0.7632	0.7632	0.7763	0.7368	0.6974
Top-5 Acc.	0.8421	0.8553	0.7895	0.8026	0.8026
Top-10 Acc.	0.9342	0.9211	0.9079	0.9079	0.9211
Recall@1	0.4178	0.4090	0.4265	0.3662	0.3882
Recall@3	0.6776	0.6919	0.7018	0.6458	0.6261
Recall@5	0.7719	0.7851	0.7390	0.7292	0.7292
Recall@10	0.8849	0.8805	0.8717	0.8465	0.8596
MAP	0.6109	0.6103	0.6171	0.5589	0.5754
MRR	0.6655	0.6526	0.6564	0.6028	0.6138

strategies, and hyperparameter choices. We mitigated these risks through fixed random seeds, iterative stratified splits, leakage-safe fitting procedures (e.g., fitting vectorizers, scalers, and binarizers on training data only), and consistent tuning protocols across models. Nevertheless, alternative configurations could lead to different outcomes. For conclusion validity, we reported multiple complementary ranking metrics (Top- k , Recall@ k , MAP, MRR). However, small differences between strong baselines should be interpreted cautiously given the sample size and residual noise in historical bug reports and commit links. We reduced noise by filtering incomplete entries and enforcing consistent multi-label encoding, but some inconsistencies likely remained. Finally, the dataset is inherently confidential. All data extraction, preprocessing, model training, and evaluation were conducted locally within ABB’s secure computing environment. No proprietary data were transmitted to external servers or cloud services. Although the dataset cannot be released, we documented the full methodology, implementation steps, and evaluation procedures and provide a replication package excluding proprietary content to support methodological reproducibility on other datasets.

7 Conclusion and Future Work

We evaluated text-only fault localization in an industrial setting using proprietary data from a large ABB Robotics system. We framed the problem as multi-label ranking over subfolders and compared three traditional models (Logistic Regression, Support Vector Machine, and Random Forest) with fine-tuned RoBERTa variants. On the held-out test set, TF-IDF-based LR and SVM consistently outperformed the transformer models, while RF became competitive when paired with simple augmentation. Sentence embeddings underperformed TF-IDF for the traditional models. The best configurations achieved Top-1 around 0.53, Top-5 up to 0.86, MAP near 0.61–0.62, and MRR around 0.66, indicating that short ranked lists can substantially narrow developers’ search space during triage. Combined with our confidentiality-friendly, fully local execution, these results support the immediate feasibility of an assistive, text-based localization capability that integrates into industrial maintenance workflows. Our study covered a single proprietary system with 660 resolved reports and 31 subfolder labels under pronounced class imbalance. While these conditions reflect common industrial constraints, replication on additional systems is required to assess transferability across domains, architectures, and reporting

practices. We provide all implementation artifacts (excluding proprietary data) as a replication package to enable reuse on comparable datasets.

We pursue three directions for future work. First, we will replicate the study across additional industrial systems and vary label granularity (file, component, and service) to understand how granularity affects accuracy and usefulness. Second, we will design hybrid, privacy-preserving models that enrich bug-report text with triage-time signals such as commit metadata, historical co-change patterns, and lightweight repository statistics, while avoiding direct ingestion of source code. Third, we will run human-in-the-loop and deployment studies to quantify effort savings (e.g., time to first correct location), calibrate confidence and abstention behavior, and integrate ranked recommendations into bug trackers and continuous integration workflows.

8 Data Availability

We provide all scripts, implementation artifacts, and supporting analysis as a replication package at <https://zenodo.org/records/18841230>. Due to confidentiality constraints, the underlying ABB Robotics bug reports and linked code-change data cannot be shared publicly. The package includes preprocessing, training, and evaluation scripts, configuration files, and step-by-step instructions to rerun the pipeline on comparable datasets. We do not distribute model checkpoints trained on proprietary data; instead, we provide hyperparameters and aggregated results to support reuse.

Acknowledgments

This work is supported by (a) the European HORIZON-KDT-JU research project MATISSE “Model-based engineering of Digital Twins for early verification and validation of Industrial Systems”, HORIZON-KDT-JU-2023-2-RIA, Proposal number: 101140216-2, KDT-232RIA00017, (b) the Sweden’s innovation agency Vinnova through the project iSecure (202301899), and (c) the Clean Energy Transition Partnership through the project “FLEXI: Human-centered AI and digital twin powered energy system integration for flexibility markets” (101069750).

References

- [1] Rui Abreu and Arjan J. C. van Gemund. 2009. A Low-Cost Approximate Minimal Hitting Set Algorithm and its Application to Model-Based Diagnosis. In *Eighth Symposium on Abstraction, Reformulation, and Approximation, SARA 2009, Lake Arrowhead, California, USA, 8-10 August 2009*, Vadim Bulitko and J. Christopher

- Beck (Eds.). AAAI. <http://www.aaai.org/ocs/index.php/SARA/SARA09/paper/view/834>
- [2] Hafiza Anisa Ahmed, Narmeen Zakaria Bawany, and Jawwad Ahmed Shamsi. 2021. CaPBug-A Framework for Automatic Bug Categorization and Prioritization Using NLP and Machine Learning Algorithms. *IEEE Access* 9 (2021), 50496–50512. doi:10.1109/ACCESS.2021.3069248
 - [3] Renato Andrade, César Teixeira, Nuno Laranjeiro, and Marco Vieira. [n. d.]. An Empirical Study on the Classification of Bug Reports with Machine Learning. arXiv. Accessed: April 27, 2025. [Online]. Available: <https://arxiv.org/abs/2503.00660>.
 - [4] Luciano C. Ascari, Lucilia Yoshie Araki, Aurora T. R. Pozo, and Silvia R. Vergilio. 2009. Exploring machine learning techniques for fault localization. In *10th Latin American Test Workshop, LATW 2009, Rio de Janeiro, Brazil, March 2-5, 2009*. IEEE, 1–6. doi:10.1109/LATW.2009.4813783
 - [5] Lionel C. Briand, Yvan Labiche, and Xuetao Liu. 2007. Using Machine Learning to Support Debugging with Tarantula. In *ISSRE 2007, The 18th IEEE International Symposium on Software Reliability, Trollhättan, Sweden, 5-9 November 2007*. IEEE Computer Society, 137–146. doi:10.1109/ISSRE.2007.31
 - [6] Ing-Xiang Chen, Cheng-Zen Yang, Ting-Kun Lu, and Hojun Jaygarl. 2008. Implicit Social Network Model for Predicting and Tracking the Location of Faults. In *2008 32nd Annual IEEE International Computer Software and Applications Conference*. 136–143. doi:10.1109/COMPSAC.2008.162
 - [7] Holger Cleve and Andreas Zeller. 2005. Locating causes of program failures. In *27th International Conference on Software Engineering (ICSE 2005), 15-21 May 2005, St. Louis, Missouri, USA*, Gruia-Catalin Roman, William G. Griswold, and Bashar Nuseibeh (Eds.). ACM, 342–351. doi:10.1145/1062455.1062522
 - [8] Vidroha Debroy and W. Eric Wong. 2009. Insights on Fault Interference for Programs with Multiple Bugs. In *ISSRE 2009, 20th International Symposium on Software Reliability Engineering, Mysuru, Karnataka, India, 16-19 November 2009*. IEEE Computer Society, 165–174. doi:10.1109/ISSRE.2009.14
 - [9] Gerhard Friedrich, Markus Stumptner, and Franz Wotawa. 1999. Model-Based Diagnosis of Hardware Designs. *Artif. Intell.* 111, 1-2 (1999), 3–39. doi:10.1016/S0004-3702(99)00034-X
 - [10] Meng Gao, Pengyu Li, Congcong Chen, and Yunsong Jiang. 2018. Research on Software Multiple Fault Localization Method Based on Machine Learning. *MATEC Web of Conferences* 232 (Nov. 2018). doi:10.1051/mateconf/201823201060
 - [11] Reza Gharibi, Amir Hossein Rasekh, Mohammad Hadi Sadreddini, and Seyed Mostafa Fakhrahmad. 2018. Leveraging textual properties of bug reports to localize relevant source files. *Information Processing and Management* 54, 6 (2018), 1058–1076. doi:10.1016/j.ipm.2018.07.004
 - [12] Thomas Hirsch and Birgit Hofer. 2022. Using textual bug reports to predict the fault category of software bugs. *Array* 15 (Sept. 2022). doi:10.1016/j.array.2022.100189
 - [13] Maliheh Izadi. 2022. Catlss: An Intelligent Tool for Categorizing Issues Reports using Transformers. In *2022 IEEE/ACM 1st International Workshop on Natural Language-Based Software Engineering (NLBSE)*. 44–47. doi:10.1145/3528588.3528662
 - [14] Sungmin Kang, Gabin An, and Shin Yoo. 2024. A Quantitative and Qualitative Evaluation of LLM-Based Explainable Fault Localization. *Proc. ACM Softw. Eng.* 1, Article 64 (Jul. 2024), 23 pages. Issue FSE. doi:10.1145/3660771
 - [15] Omer Koksal and Bedir Tekinerdogan. 2022. Automated Classification of Unstructured Bilingual Software Bug Reports: An Industrial Case Study Research. *Applied Sciences* 12, 1 (2022). doi:10.3390/app12010338
 - [16] Zhengliang Li, Zhiwei Jiang, Xiang Chen, Kaibo Cao, and Qing Gu. 2020. Laprob: A Label propagation-Based software bug localization method. *Information and Software Technology* 130 (Oct. 2020). doi:10.1016/j.infsof.2020.106410
 - [17] Ben Liblit, Mayur Naik, Alice X. Zheng, Alex Aiken, and Michael I. Jordan. 2005. Scalable statistical bug isolation. In *Proceedings of the ACM SIGPLAN 2005 Conference on Programming Language Design and Implementation, Chicago, IL, USA, June 12-15, 2005*, Vivek Sarkar and Mary W. Hall (Eds.). ACM, 15–26. doi:10.1145/1065010.1065014
 - [18] Chao Liu, Long Fei, Xifeng Yan, Jiawei Han, and Samuel P. Midkiff. 2006. Statistical Debugging: A Hypothesis Testing-Based Approach. *IEEE Trans. Software Eng.* 32, 10 (2006), 831–848. doi:10.1109/TSE.2006.105
 - [19] Edward Ogheneovo. 2014. On the Relationship between Software Complexity and Maintenance Costs. *Journal of Computer and Communications* 02 (Jan. 2014), 1–16. doi:10.4236/jcc.2014.214001
 - [20] Priyan Vaithilingam, Tianyi Zhang, and Elena L. Glassman. 2022. Expectation vs. experience: Evaluating the usability of code generation tools powered by large language models. In *Chi conference on human factors in computing systems extended abstracts*. 1–7.
 - [21] Neil Walkinshaw and Leandro Minku. 2018. Are 20% of files responsible for 80% of defects?. In *Proceedings of the 12th ACM/IEEE International Symposium on Empirical Software Engineering and Measurement (Oulu, Finland) (ESEM '18)*. Association for Computing Machinery, New York, NY, USA, Article 2, 10 pages. doi:10.1145/3239235.3239244
 - [22] Junjie Wang, Yuchao Huang, Chunyang Chen, Zhe Liu, Song Wang, and Qing Wang. 2024. Software testing with large language models: Survey, landscape, and vision. *IEEE Transactions on Software Engineering* 50, 4 (2024), 911–936.
 - [23] W. Eric Wong, Vidroha Debroy, and Dianxiang Xu. 2012. Towards Better Fault Localization: A Crosstab-Based Statistical Approach. *IEEE Trans. Syst. Man Cybern. Part C* 42, 3 (2012), 378–396. doi:10.1109/TSMCC.2011.2118751
 - [24] W. Eric Wong, Ruizhi Gao, Yihao Li, Rui Abreu, and Franz Wotawa. 2016. A Survey on Software Fault Localization. *IEEE Transactions on Software Engineering* 42, 8 (Aug. 2016), 707–740. doi:10.1109/TSE.2016.2521368
 - [25] W. Eric Wong and Yu Qi. 2009. Bp Neural Network-Based Effective Fault Localization. *Int. J. Softw. Eng. Knowl. Eng.* 19, 4 (2009), 573–597. doi:10.1142/S021819400900426X
 - [26] Chunqiu Steven Xia and Lingming Zhang. 2023. Conversational automated program repair. *arXiv preprint arXiv:2301.13246* (2023).
 - [27] Chuyang Xu, Zhongxin Liu, Xiaoxue Ren, Gehao Zhang, Ming Liang, and David Lo. 2025. FlexFL: Flexible and Effective Fault Localization With Open-Source Large Language Models. *IEEE Transactions on Software Engineering* 51, 5 (2025), 1455–1471. doi:10.1109/TSE.2025.3553363
 - [28] Aidan Z. H. Yang, Claire Le Goues, Ruben Martins, and Vincent Hellendoorn. 2024. Large Language Models for Test-Free Fault Localization. In *Proceedings of the IEEE/ACM 46th International Conference on Software Engineering (Lisbon, Portugal) (ICSE '24)*. Association for Computing Machinery, New York, NY, USA, Article 17, 12 pages. doi:10.1145/3597503.3623342
 - [29] Haiyang Yang, Yulu Zhou, Tian Liang, and Li Kuang. 2025. ChatDL: An LLM-Based Defect Localization Approach for Software in IIoT Flexible Manufacturing. *IEEE Internet of Things Journal* 12, 16 (2025), 32333–32343. doi:10.1109/JIOT.2025.3531512
 - [30] Andreas Zeller. 2009. *Why Programs Fail - A Guide to Systematic Debugging, 2nd Edition*. Academic Press. <http://store.elsevier.com/product.jsp?isbn=9780123745156&pagename=search>
 - [31] Andreas Zeller and Ralf Hildebrandt. 2002. Simplifying and Isolating Failure-Inducing Input. *IEEE Trans. Software Eng.* 28, 2 (2002), 183–200. doi:10.1109/32.988498
 - [32] Jialu Zhang, José Cambronero, Sumit Gulwani, Vu Le, Ruzica Piskac, Gustavo Soares, and Gust Verbruggen. 2022. Repairing bugs in python assignments using large language models. *arXiv preprint arXiv:2209.14876* (2022).
 - [33] Hu Zhengdong, Polpini Jantima, and Gangarn Somprasertsri. 2025. A Comparative Study of Multilabel Classification Techniques for Analyzing Bug Report Dependencies. *Journal of Information Systems Engineering and Management* 10 (2025). doi:10.52783/jisem.v10i26s.4253