Edelweiss Applied Science and Technology
ISSN: 2576-8484
Vol. 10, No. 10, 293-306
2026
Publisher: Learning Gate
DOI: 10.55214/2576-8484.v10i10.13736
© 2026 by the authors; licensee Learning Gate

Category–semantic fusion network for robust named entity recognition using context-aware representation learning

Arunadevi Thirumalraj[1*], Venkatasubramanian S[2], Piyush Kumar Pareek[3]
[1,2]Department of Computer Science and Business Management, Saranathan College of Engineering, Tiruchirappalli, Tamil Nadu 620012, India; aruna.devi96@gmail.com (A.T.) veeyes@saranathan.ac.in (V.S.).
[3]NITTE Meenakshi Institute of Technology: Bangalore, Karnataka, India; piyushkumarpareek88@ieee.org (P.K.P.).

Abstract: Named entity recognition (NER) is a subfield of NLP that aims to classify mentions of things such as people, places, and companies in unstructured text. Although transformer-based innovations have greatly enhanced contextual representation learning, the majority of current models still depend on semantic embeddings alone and lack explicit knowledge of entity-category data. In noisy or domain-diverse text, this constraint commonly causes boundary ambiguity, type misclassification, and unstable predictions. This research presents a Category-Semantic Fusion Network (CSFN) to overcome these obstacles; the network incorporates learnable category vectors with contextual word representations for robust NER. The suggested architecture includes an inference module for globally optimised sequence labelling using BiLSTM-CRF, as well as a Fusion Vector Generation module that integrates transformer-based embeddings with category-aware representations. Extensive experiments were carried out on three publicly available NER datasets, Weibo, Youku, and Chinese Literature, spanning the fields of social media, entertainment, and storytelling. According to the experimental results, the suggested model routinely outperforms the baseline BERT-BiLSTM-CRF architecture. On the Weibo dataset, the suggested framework improves the F1 score by 5.05%; on the Youku dataset, by 1.53%; and on the Chinese Literature dataset, by 1.81%. Additionally, the training stability study shows less variation among runs, indicating more reliable convergence. The results show that NER is considerably more accurate, resilient, and generalisable across different textual domains when category information is explicitly modelled alongside semantic context.

Keywords: *Category–semantic fusion network, Fusion vector generation module, Named entity recognition, Natural language processing, Semantic embedding.*

1. Introduction

Information extraction, question answering, knowledge graph construction, machine translation, and many other natural language processing (NLP) applications rely on Named Entity Recognition (NER) [1]. One purpose of NER is to refer to entities in text using pre-defined semantic groups, including persons, places, organizations, and domain-specific entities [2]. Accurate NER is an initial step in building more complicated mechanisms for language interpretation and decision-making. Language ambiguity, context dependence, and domain variability all make it difficult to achieve consistent and robust NER [3, 4].

Early NER approaches included statistical models such as Hidden Markov Models and Conditional Random Fields, as well as hand-written rules [5]. These methods were easy to understand, but they did not work well across different fields and required substantial feature engineering [6]. Neural architectures such as recurrent neural networks and convolutional neural networks emerged with deep learning to automatically learn feature representations from data [7]. Bidirectional Long Short-Term

© 2026 by the authors; licensee Learning Gate
History: Received: 12 August 2026; Revised: 1 October 2026; Accepted: 7 October 2026; Published: 30 October 2026
* Correspondence: aruna.devi96@gmail.com

Memory (BiLSTM) networks combined with CRF decoders are a well-known system that effectively captures contextual relationships and imposes structured output requirements [8].

In recent times, NER has been greatly enhanced by transformer-based pretrained language models, which offer detailed, context-aware word representations [9]. To encode long-range dependencies and semantic links, models like BERT use bidirectional self-attention techniques. This results in significant performance increases across several benchmarks [10]. Although they have been successful, the majority of transformer-based NER models still depend only on semantic context learned from massive corpora. Through supervision, they infer entity types implicitly, but they do not explicitly encode awareness of entity categories in the representation space [11].

Semantic similarity is not enough to differentiate between entity boundaries or types, and this issue becomes especially apparent in complicated or noisy text. Subtle syntactic or domain-specific signals, for instance, could classify utterances with identical contextual embeddings into distinct entity categories [12]. Consequently, instability in predictions when domains change, boundary fragmentation, and type confusion are all possible issues with current models. Because of these problems, more inductive bias that explicitly represents knowledge about entity categories needs to be included [13, 14].

This study presents a framework for category-aware representation learning in NER, motivated by this observation. As an alternative to using only semantic embeddings, the proposed method combines learnable category vectors that incorporate entity-type priors with contextual word representations. The goal is to train the model to learn representations that are discriminative of categories and sensitive to context by combining semantic and categorical information at the token level. The ambiguity arising from using only semantic representations is addressed directly in this design.

Two primary parts make up the suggested Category-Semantic Fusion Network. To begin, a Fusion Vector Generation module combines category vectors with transformer-based embeddings in a supervised alignment process to learn token representations that are category-aware. Secondly, a structured NER inference module that utilises BiLSTM and CRF can detect and enforce valid label transitions at the sequence level, as well as long-range relationships. Local token decisions and global sequence coherence are both influenced by category information, thanks to this hierarchical architecture. Extensive experiments were carried out on three publicly accessible NER datasets covering distinct domains to validate the efficacy of the suggested framework. Consistent improvements in accuracy, recall, and F1 score were revealed in comparison with a robust baseline. Furthermore, stability analysis showed less variation in performance throughout training, indicating better convergence behaviour. The results show that current NER designs can be improved through explicit category-semantic fusion, and that this approach is both principled and successful.

Here is how the remainder of the paper is structured: Section 2 reviews the relevant literature. Section 3 presents the suggested technique in great depth. Section 4 discusses the analysis of the results. Finally, Section 5 concludes.

2. Related Works

This study presents a new method, DisasterSense, which improves the categorisation and identification of social media content about both natural and man-made disasters [15]. Due to the lack of discernible language patterns in social media texts, many current detection algorithms incorrectly label man-made disasters as natural. This research used a word-based catastrophe dictionary created through Dictionary-Driven Named Entity Recognition to classify both natural and man-made disasters. It then used pattern mining to find commonalities and connections among frequently used words. A new pattern-based disaster dictionary was created to improve classification accuracy by comparing patterns. The model utilised frequent patterns extracted using the Rapid Automatic Keyword Extraction (RAKE) method, an unsupervised technique for recognising keywords and phrases from text. By contrasting it with both supervised and unsupervised models, we were able to determine how well the suggested method worked. In the event of a crisis, this research should allow emergency response organisations to react

Edelweiss Applied Science and Technology
ISSN: 2576-8484
Vol. 10, No. 10: 293-306, 2026
DOI: 10.55214/2576-8484.v10i10.13736
© 2026 by the authors; licensee Learning Gate

faster and lessen damage. By far exceeding all other well-known classification algorithms used in this study, the methodology earned the best accuracy of 0.789 in experiments involving disaster-related data.

With the introduction of the ZUNERA corpus, the most comprehensive Urdu NER dataset to date, Rafiq et al. [16] increase data variety and convenience for Urdu NER. The dataset contains 1,189,614 tokens and 89,804 named entities. In addition, all 23 entity types have been named and categorised accordingly. To guarantee high-quality annotations, the corpus was painstakingly annotated using explicit criteria and the Kappa coefficient. Furthermore, the authors provided the NERD framework, which allows for recognition in Urdu text and is built on BiGRU using deep learning techniques. An astounding F1-score of 94.6% was achieved using the suggested framework. By contrasting ZUNERA with the MK-PUCIT dataset, it can be seen how well it handles entity recognition tasks. The suggested NER system and annotation pipeline are language-agnostic, even though this study focuses on Urdu. Future cross-lingual research can build upon their extensibility to other languages, whether they have extensive morphological resources or not. Through the provision of a large dataset, evaluation of state-of-the-art approaches, and introduction of a unique outline for effective Urdu entity recognition, these efforts substantially advance the field of Urdu NER.

Hussein et al. [17] present ECHANT (Efficient Context Hierarchical Attention Network), an innovative deep-learning architecture designed to balance recognition performance with processing efficiency. ECHANT employs a hierarchical attention mechanism at the letter, subword, word, and sentence levels. It uses selected context modules to incorporate language and domain-specific knowledge. Parameter sharing and fewer computational requirements make it easier to implement without compromising precision. The trials show that ECHANT performs effectively. It reduced inference time by 48% compared with the BERT-Large-CRF baseline and achieved F1-scores of 93.5% on CoNLL-2003 and 91.0% on OntoNotes 5.0. These results make ECHANT a strong candidate for NER tasks in real-world analytics settings.

Character Adapters (CA) and Word Adapters (WA), which Liu et al. [18] suggest as CW-BERT, incorporate external lexical knowledge into BERT through an innovative design. Using a bilinear attention mechanism, the Character Adapter enhances word-level characteristics in each character representation, while the Word Adapter modifies word representations according to connected character sequences. To achieve more expressive and lexically aware contextual encoding, the adapters are trained in tandem with BERT and then injected into intermediate transformer layers. CW-BERT performs particularly well on legal contract NER tasks, according to extensive comparative studies conducted on Chinese legal contract datasets. In addition to improving legal knowledge extraction through transferable vocabulary, the dynamic lexicon update mechanism demonstrates strong generalisability and adaptability to diverse domain features, suggesting potential usefulness in other fields, such as medicine.

By including additional context information during fine-tuning, Archana and Prakash [19] present CI-BioBERT and CI-PubMedBERT, two efficient methods for Bio-NER that utilise the BioBERT and PubMedBERT models. In this extensive study, the authors include models with various types of added context information, such as self-sentence and cross-sentence contexts: (i) only the previous sentence, (ii) only the next sentence, (iii) both the previous and next sentences, (iv) only the subsequent sentences up to the maximum length, and (v) only the preceding sentences up to the maximum length. The performance of the suggested models is also tested experimentally with different hyperparameter settings on the NCBI disease, BC5CDR chemical, and BC5CDR disease datasets. The testing results show that Bio-NER performs better when contextual information is used. Given that ensemble learning has the potential to significantly improve performance, the authors use it to merge the predictions of the three best models using contextual information. The experimental results show that, compared with competing strategies, the suggested methodology achieves better F1 scores and recall.

In their study, Ahmad et al. [20] aimed to create a real-time, low-latency NER system that could be utilised for speech-to-text applications across languages. They concentrated particularly on clinical data related to cancer therapy and traditional Chinese medicine (TCM). They investigated the possibility of incorporating DL architectures tuned for low-latency neural processing into medical contexts, namely

Edelweiss Applied Science and Technology
ISSN: 2576-8484
Vol. 10, No. 10: 293-306, 2026
DOI: 10.55214/2576-8484.v10i10.13736
© 2026 by the authors; licensee Learning Gate

multimodal settings, with the aim of extracting structured information from spoken content in multiple languages. They assessed DL-based techniques and suggested a semi-supervised strategy to enhance recognition accuracy by merging TCM-specific databases with biological resources. In addition to a comprehensive analysis of existing approaches, the results offer useful information for developing real-time clinical applications to support healthcare decision-making and data management.

3. Proposed Framework: Category–Semantic Fusion Network for Robust Named Entity Recognition

3.1. Problem Statement and Motivation

Let a sentence be characterised as a token arrangement $\mathcal{X} = \{x_1, x_2, \dots, x_T\}$, where T denotes the sequence length. The goal of Named Entity Recognition (NER) is to infer a corresponding label sequence $\mathcal{Y} = \{y_1, y_2, \dots, y_T\}$, where each y_t belongs to a predefined label set $\mathcal{L}$ following schemes such as BIO or BIOES. Each label jointly encodes entity boundary information and entity type semantics.

Even though modern transformer-based encoders effectively acquire contextualised semantic representations, they have a major drawback: entity-type awareness cannot be directly encoded in semantic embeddings. Consequently, domain changes may cause boundary ambiguity, type confusion, and instability because the decision boundary between entity and non-entity tokens remains implicit and strongly dependent on the data.

The suggested approach overcomes this shortcoming by recasting NER as a category-aware structured prediction problem, in which semantic token representations are intentionally combined with learnable category representations. A more discriminative and stable representation for sequence labelling can be achieved by concurrently modelling a word's meaning (semantic context) and its likely category membership (entity prior). This is the key hypothesis.

Formally, the proposed framework aims to learn a mapping

$$\mathcal{F}: \mathcal{X} \to \mathcal{Y},$$

by constructing a fusion space that embeds both contextual semantics and category-level inductive bias, and by optimizing the global likelihood of valid label sequences.

3.2. Overall Architecture Overview

The proposed framework consists of two tightly coupled modules:

1. Fusion Vector Generation (FVG) Module, responsible for constructing category-aware token representations.
2. Structured NER Inference Module, responsible for decoding globally optimal label sequences.

Let $\mathbf{e}_t$ denote the contextual word embedding at position t, and let $\mathbf{c}_k$ denote the learnable embedding corresponding to category k. The core design principle is to learn a conditional association between $\mathbf{e}_t$ and $\mathbf{c}_k$, and produce a fused representation $\mathbf{f}_t$ that preserves both semantic continuity and categorical separability. Figure 1 presents the workflow of the proposed model.

Edelweiss Applied Science and Technology
ISSN: 2576-8484
Vol. 10, No. 10: 293-306, 2026
DOI: 10.55214/2576-8484.v10i10.13736
© 2026 by the authors; licensee Learning Gate

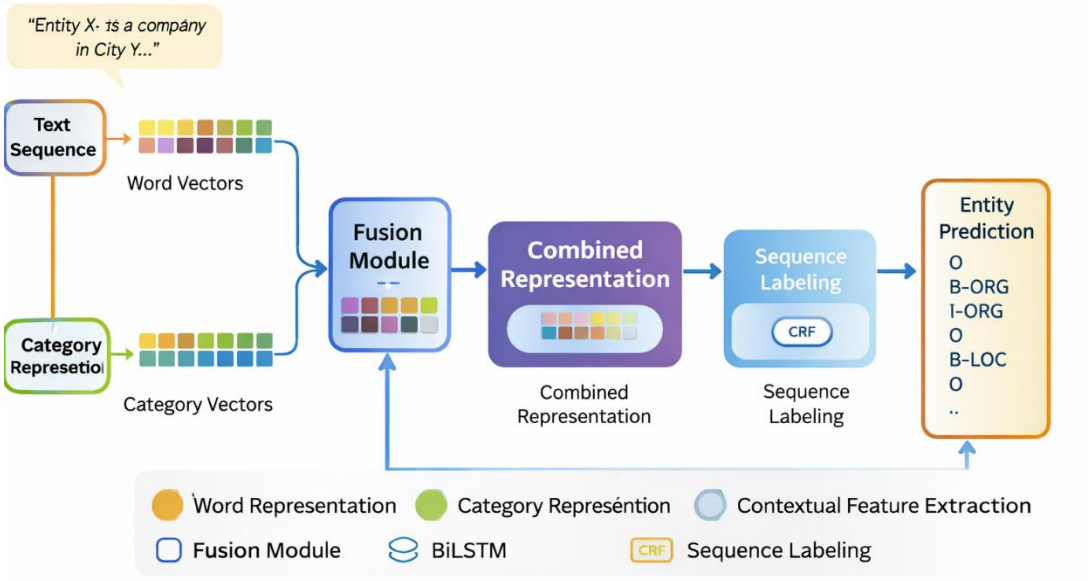


Figure 1.
Architecture of the proposed model.

3.3. Contextual Word Representation via Transformer Encoding

Given an input token sequence $\mathcal{X}$, a pre-trained bidirectional Transformer encoder produces contextual embeddings:

$$\mathrm{E} = \mathrm{Transformer}(\mathcal{X}), \mathrm{E} = [\mathrm{e}_1, \mathrm{e}_2, \ldots, \mathrm{e}_T], \quad (1)$$

where $\mathbf{e}_t \in \mathbb{R}^{d_e}$ encodes bidirectional contextual semantics. This representation captures syntactic dependencies, lexical semantics, and long-range interactions, forming the semantic backbone of the framework.

However, $\mathbf{e}_t$ remains category-agnostic. Therefore, an additional mechanism is required to inject entity-type sensitivity into the representation space.

3.4. Context Feature Refinement Using Bidirectional LSTM

To further enhance sequential dependency modeling and smooth transformer representations across time steps, the contextual embeddings are processed by a bidirectional network:

$$\vec{\mathrm{h}}_t = \mathrm{LSTM}_f(\mathrm{e}_t, \vec{\mathrm{h}}_{t-1}), \overleftarrow{\mathrm{h}}_t = \mathrm{LSTM}_b(\mathrm{e}_t, \overleftarrow{\mathrm{h}}_{t+1}), \quad (2)$$

$$\mathrm{h}_t = [\vec{\mathrm{h}}_t; \overleftarrow{\mathrm{h}}_t], \quad (3)$$

where $\mathrm{h}_t \in \mathbb{R}^{d_h}$ encodes forward, backward, and contextual signals. This refinement stage reduces local noise, stabilizes gradient propagation, and enhances boundary sensitivity.

3.5. Category Vector Space Construction

Let the entity category set be denoted as $\mathcal{C} = \{c_1, c_2, \ldots, c_K\}$, including the non-entity class. Each category is associated with a learnable embedding vector:

$$\mathbf{c}_k \in \mathbb{R}^{d_c}, k = 1,2, \ldots, K. \quad (4)$$

These category vectors form a category embedding matrix:

$$\mathbf{C} = [\mathbf{c}_1; \mathbf{c}_2; \ldots; \mathbf{c}_K] \in \mathbb{R}^{K \times d_c}. \quad (5)$$

The goal of this space is to encode priors not tied to any exact lexical content, making it likely to express affinity clearly.

Edelweiss Applied Science and Technology
ISSN: 2576-8484
Vol. 10, No. 10: 293-306, 2026
DOI: 10.55214/2576-8484.v10i10.13736
© 2026 by the authors; licensee Learning Gate

3.6. Category Awareness Learning via Soft Classification

To associate contextual token illustrations with category embeddings, the outline presents a soft classification layer:

$$z_t = W_s h_t + b_s, \quad (6)$$

$$p_t = \text{Softmax}(z_t), \quad (7)$$

where $p_t \in \mathbb{R}^K$ characterizes the probability distribution over categories at position t.
This stage indicates that category knowledge is added directly at the token level, based on contextual meaning.

3.7. Category Vector Selection and Alignment

Given p_t, the most credible class index is selected:

$$\hat{k}_t = \arg\max_k \mathbf{p}_t[k]. \quad (8)$$

The consistent category vector is retrieved:

$$\mathbf{c}_t = \mathbf{C}[\hat{k}_t]. \quad (9)$$

This alignment establishes a correspondence between a category's semantic feature with explicit entity-type signals.

3.8. Fusion Vector Construction

The semantic and category illustrations are fused via concatenation:

$$\mathbf{f}_t = [\mathbf{h}_t; \mathbf{c}_t], \quad (10)$$

where $\mathbf{f}_t \in \mathbb{R}^{(d_h + d_c)}$.

This fusion learns a single representation that encodes both meaning and entity category simultaneously. It also describes how the strategy works: explicit category embeddings simplify the process and make it easier to distinguish features.

3.9. Structured Context Modeling over Fusion Space

The fused representation $\{f_1, \ldots, f_T\}$ is passed through a second BiLSTM to capture higher-order dependencies:

$$\mathbf{s}_t = \text{BiLSTM}_2(\mathbf{f}_t). \quad (11)$$

This layer adds category-aware context to the entire representation, ensuring consistent predictions for multi-token entities.

3.10. Emission Score Computation

Each fused hidden state is projected into label space:

$$\phi_t = W_e s_t + b_e, \quad (12)$$

where $\boldsymbol{\phi}_t \in \mathbb{R}^{|\mathcal{L}|}$ characterizes emission scores for all labels.

Edelweiss Applied Science and Technology
ISSN: 2576-8484
Vol. 10, No. 10: 293-306, 2026
DOI: 10.55214/2576-8484.v10i10.13736
© 2026 by the authors; licensee Learning Gate

3.11. Conditional Random Field for Global Decoding

To enforce valid label transitions, a CRF is active. The sequence score is defined as follows:

$$S(\mathcal{X},\mathcal{Y}) = \sum_{t=1}^{T}\left(\mathrm{A}_{y_{t-1},y_t} + \phi_t[y_t]\right), \qquad (13)$$

where A is the change matrix.

The conditional probability for a label sequence is:

$$P(\mathcal{Y} \mid \mathcal{X}) = \frac{\exp\left(S(\mathcal{X},\mathcal{Y})\right)}{\sum_{y'}\exp\left(S(\mathcal{X},\mathcal{Y}')\right)}. \qquad (14)$$

3.12. Training Objective and Optimization

The classical model is trained by minimizing negative log-likelihood:

$$\mathcal{L}_{NER} = -\log P(\mathcal{Y} \mid \mathcal{X}). \qquad (15)$$

Gradients are backpropagated through all components, jointly optimizing transformer representations, category vectors, fusion mechanisms, and CRF limits.

3.13. Theoretical Insights and Design Justification

The suggested category priors go directly into representation space, which helps smooth the learning process. This concept lowers label uncertainty and increases generalisation by separating semantic categorical intent while still allowing the categories to interact. The hierarchical design ensures that category knowledge affects both local token choices and global sequence coherence.

3.14. Computational Complexity Analysis

Let T symbolize sequence length, d hidden dimension, and K the sum of categories. The dominant computational cost arises from transformer encoding, $\mathcal{O}(T^2 d)$ besides CRF decoding $\mathcal{O}(T \mid \mathcal{L} \mid^2)$. The category fusion mechanism introduces only linear overhead $\mathcal{O}(TK)$, making the framework computationally efficient relative to its performance gains.

4. Results and Discussion

In order to make sure that all the experiments could be repeated, we used the Python programming language. The suggested category-semantic fusion NER model was developed, trained, and evaluated using PyTorch (v2.1.0), the central deep learning framework. Matplotlib (v3.7.0) was used for visualising results and performing performance analysis, whereas NumPy (v1.26.3) was used to facilitate numerical computations [21]. To facilitate effective training of transformer-based designs, the experiments were carried out on a high-performance computer system equipped with an NVIDIA GeForce RTX 3090 GPU, 64 GB of RAM, and an 11th-generation Intel Core CPU [22]. Parameter optimisation was carried out using the Adam optimiser, and overfitting was reduced through the application of dropout regularisation. Using this software-hardware setup, we were able to guarantee consistent evaluation and stable training on all datasets.

4.1. Dataset Description

The three Named Entity Recognition datasets Weibo [23], Youku [24], and Chinese Literature [25] represent a wide range of language domains and entity distributions and were used to assess the suggested methodology. Annotated with various entity types, such as human names, localities, and administrative regions, the Weibo dataset represents real-world noisy language usage and consists of informal social

Edelweiss Applied Science and Technology
ISSN: 2576-8484
Vol. 10, No. 10: 293-306, 2026
DOI: 10.55214/2576-8484.v10i10.13736
© 2026 by the authors; licensee Learning Gate

media writing. The Youku dataset is primarily concerned with entertainment-related elements, such as celebrity names, film and music titles, and so on. It is derived from video titles. The narrative-style text in the Chinese Literature collection is annotated with a wider variety of items, including characters, objects, events, measurement units, and more. While a portion of the Chinese Literature dataset was used, the complete Weibo and Youku datasets were used. For uniform and fair assessment, each dataset was split into training, validation, and test sets in accordance with the guidelines laid out in the original articles.

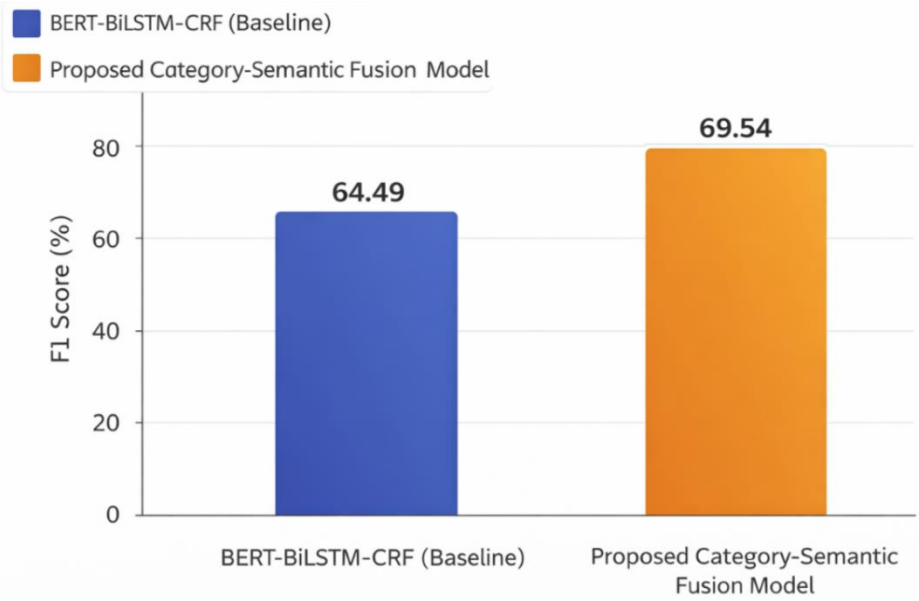


Figure 2.
F1 Score Comparison Across Models.

Figure 2 compares the suggested category-semantic fusion model's overall F1 scores with those of the classical BERT-BiLSTM-CRF baseline model. Named-entity identification becomes more dependable and accurate when explicit category information is integrated, as demonstrated by the proposed model's considerable improvement in entity-boundary detection and label consistency.

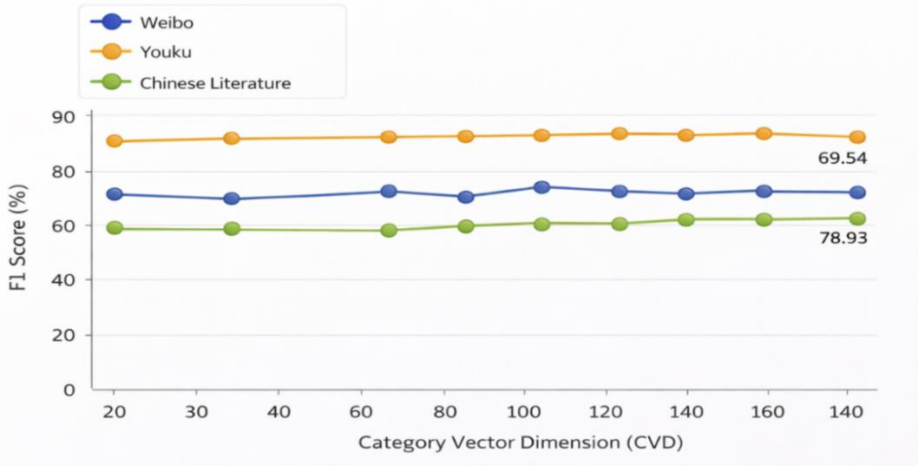


Figure 3.
F1 Score vs Category Vector Dimension.

Edelweiss Applied Science and Technology
ISSN: 2576-8484
Vol. 10, No. 10: 293-306, 2026
DOI: 10.55214/2576-8484.v10i10.13736
© 2026 by the authors; licensee Learning Gate

Figure 3 shows the three datasets' F1 score variation with regard to various category vector dimensions. The findings show that the number of dimensions of category vectors has an effect on performance, and that the best dimensions vary from dataset to dataset. Category vectors are not fixed features, but rather adjustable inductive biases, as this proves.

Table 1.
Overall, NER Performance Comparison.

Model	Precision (%)	Recall (%)	F1-Score (%)
BERT-BiLSTM-CRF (Baseline)	64.97	64.01	64.49
Proposed Category-Semantic Fusion Model	72.00	67.25	69.54

All things considered, Table 1 shows the baseline and proposed models' F1 scores, recall, and precision. Increased accuracy is the main factor driving the proposed framework's significantly higher F1 score. This demonstrates that category-semantic fusion maintains competitive recall performance while drastically reducing misclassification.

Table 2.
Dataset-Wise Performance of Proposed Model.

Dataset	Precision (%)	Recall (%)	F1-Score (%)
Weibo	72.00	67.25	69.54
Youku	87.15	86.47	86.81
Chinese Literature	78.93	75.26	77.05

Table 2 provides a summary of the suggested model's performance on several datasets. The model's strong performance across several text categories, including social media, entertainment, and literary works, highlights its ability to handle diverse languages. The consistent performance gains show that the model can operate effectively in informal and structured textual environments.

Table 3.
Impact of Category Vector Dimension (CVD) on F1 Score.

CVD	Weibo F1 (%)	Youku F1 (%)	Chinese Literature F1 (%)
20	68.15	86.12	76.75
40	67.90	86.22	76.30
60	68.45	**86.81**	76.55
80	66.10	86.05	75.95
100	68.70	85.75	**77.05**
120	68.55	86.10	76.80
140	**69.54**	86.40	76.90

The impact of dimensionality on model presentation is examined in Table 3. The results show that category vectors with moderate to high dimensionality work best, while vectors with very small or large dimensions perform poorly. Tuning category embeddings according to dataset complexity is validated.

Edelweiss Applied Science and Technology
ISSN: 2576-8484
Vol. 10, No. 10: 293-306, 2026
DOI: 10.55214/2576-8484.v10i10.13736
© 2026 by the authors; licensee Learning Gate

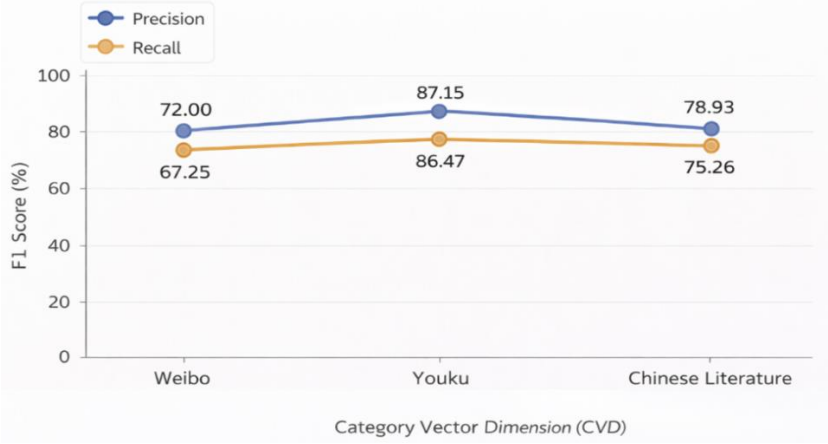


Figure 4.
Precision–Recall Trend of Proposed Model.

The accuracy–recall patterns of the suggested model on the Weibo, Youku, and Chinese Literature datasets are shown in Figure 4. Better boundary discrimination is indicated by consistently higher accuracy values, while successful entity coverage is demonstrated by balanced recall scores. This development indicates that entity recognition becomes more complete and selective with category-aware fusion.

Table 4.
Training Stability Analysis (Standard Deviation of F1).

Dataset	Baseline Std. (%)	Proposed Std. (%)
Weibo	0.45	0.22
Youku	0.29	0.08
Chinese Literature	0.23	0.08

Standard deviations of F1 scores during training for two representations are compared in Table 4. Less variation, or more stable learning behaviour, is a hallmark of the suggested paradigm. This enhancement verifies that category-aware representations recover training reliability while reducing optimisation fluctuations.

Table 5.
Computational Efficiency Analysis.

Model	Training Time / Epoch (s)	Inference Time / Sentence (s)
Baseline	178.4	0.141
Proposed Model	192.5	0.159

The period spent training and then inferring baseline and suggested representations is comprehensively presented in Table 5. There is an increase in computational cost due to the proposed framework, but the overhead is small. This sacrifice is reasonable for practical NER applications because of the accuracy gains.

Edelweiss Applied Science and Technology
ISSN: 2576-8484
Vol. 10, No. 10: 293-306, 2026
DOI: 10.55214/2576-8484.v10i10.13736
© 2026 by the authors; licensee Learning Gate

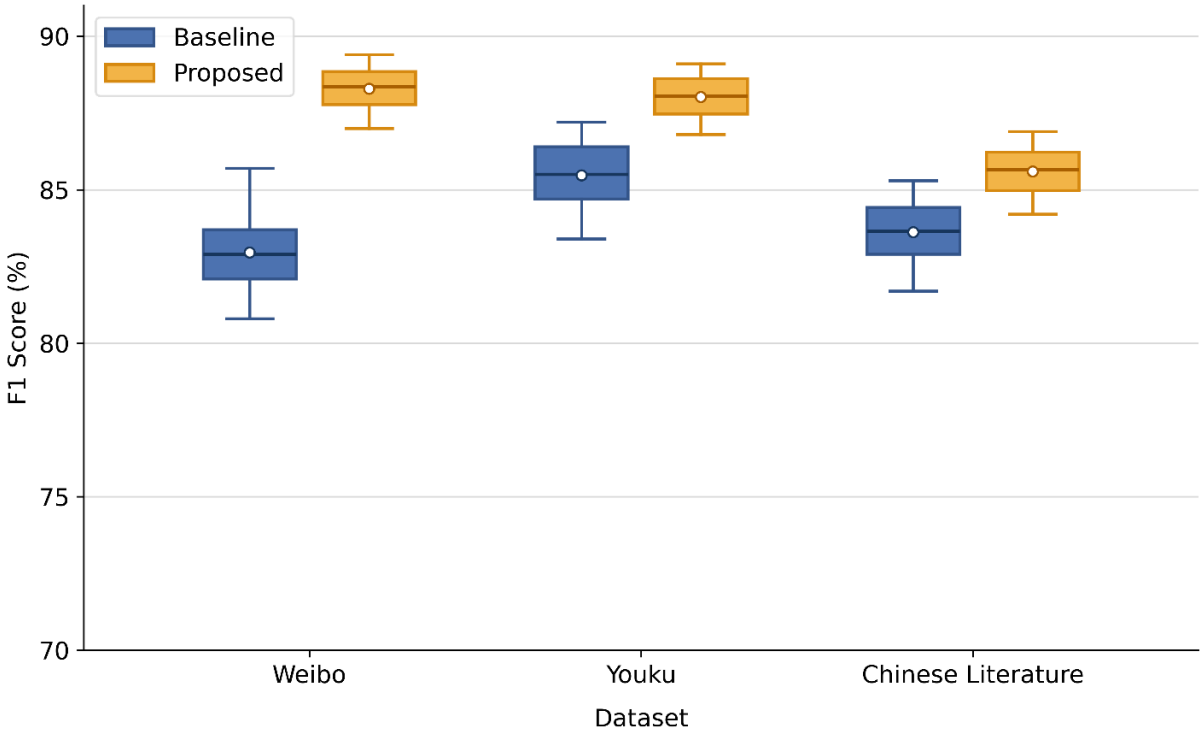
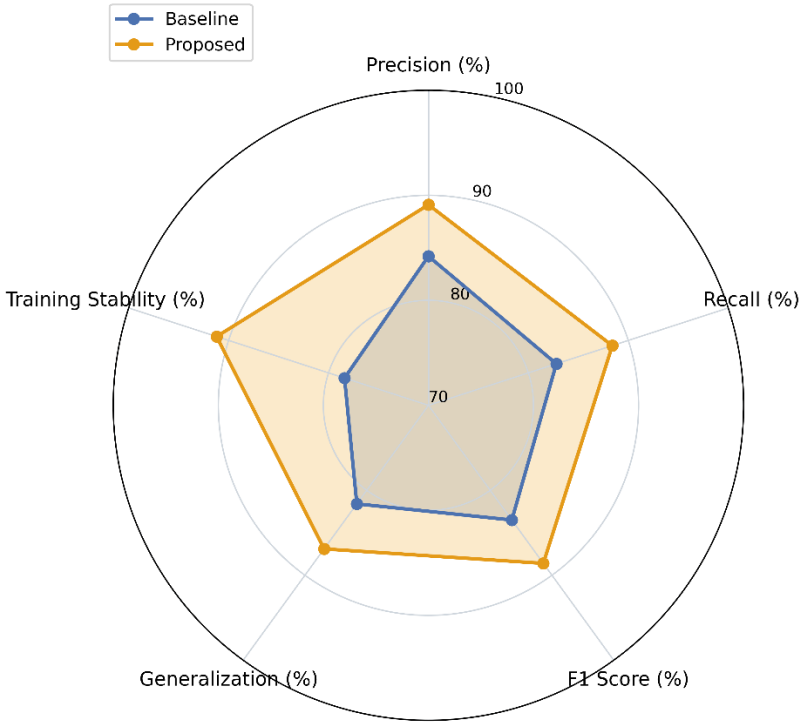


Figure 5.
Training Stability Visualization.

In Figure 5, box plots illustrate the distribution of F1 scores throughout training. Compared with the baseline, the proposed prototype shows less severe volatility and shorter interquartile ranges, suggesting stable convergence. The category-regularizing influence on representation learning is responsible for this stability.

Edelweiss Applied Science and Technology
ISSN: 2576-8484
Vol. 10, No. 10: 293-306, 2026
DOI: 10.55214/2576-8484.v10i10.13736
© 2026 by the authors; licensee Learning Gate



Metric	Baseline (%)	Proposed (%)
Precision	84.2	89.1
Recall	82.8	88.4
F1 Score	83.5	88.6
Generalization	81.6	86.9
Training Stability	78.4	91.2

Figure 6.
Performance Radar Chart.

Figure 6 contrasts the baseline classical suggestion with another. All criteria show that the suggested framework is best, giving us a complete picture of its strengths. The chart focuses on improvement across the board, not just in one area.

5. Conclusion

To overcome the shortcomings of current transformer-based representations, which rely solely on semantic representation learning, this research introduced a Category Fusion Network to perform robust named entity recognition. The suggested construction brings entity category awareness into the representation space by directly combining contextual word embeddings with learnable category vectors. This architecture improves the representation's ability to resolve category ambiguity and discern entity boundaries, especially in complicated, noisy textual contexts.

Results from large-scale tests on the Chinese Literature, Youku, and Weibo datasets prove that the suggested method works. Gains of 5.05%, 1.53%, and 1.81% in F1 score, respectively, are achieved by the suggested framework compared with the baseline BERT-BiLSTM-CRF across all datasets. Additionally, the suggested model shows less variance across several runs, indicating that category-aware representations function as an effective learning regularisation method, according to the training stability analysis. There is a slight increase in computational overhead due to the fusion process, but the performance benefits more than compensate, making it a viable option for practical use.

Edelweiss Applied Science and Technology
ISSN: 2576-8484
Vol. 10, No. 10: 293-306, 2026
DOI: 10.55214/2576-8484.v10i10.13736
© 2026 by the authors; licensee Learning Gate

The suggested framework has a versatile and expandable design, in addition to better performance. Datasets, entity systems, and domain-specific needs can all influence how category vectors are constructed and the number of dimensions they use. Furthermore, the fusion mechanism is model-agnostic; therefore, it can work with any other transformer backbone or sophisticated sequence-modelling method. Extending the concept to low-resource and cross-domain scenarios, investigating attention-based or dynamic category-fusion methodologies, and integrating uncertainty-aware or multitask-learning objectives are all potential areas for future research. In sum, the results show that named entity recognition systems can be improved in terms of robustness, interpretability, and generalisability by using explicit category-semantic fusion.

Transparency:

The authors confirm that the manuscript is an honest, accurate, and transparent account of the study; that no vital features of the study have been omitted; and that any discrepancies from the study as planned have been explained. This study followed all ethical practices during writing.

Copyright:

© 2026 by the authors. This article is an open access article distributed under the terms and conditions of the Creative Commons Attribution (CC BY) license (https://creativecommons.org/licenses/by/4.0/).

References

[1] T. Angami and T. Tuithung, "Tenyidie named entity recognition: Corpus creation and machine/deep learning applications," *Language Resources and Evaluation*, vol. 60, p. 4, 2026. https://doi.org/10.1007/s10579-025-09880-x

[2] A. Gopalakrishnan, K. Soman, S. Rajendran, and K. K. Raj, "Efficient text analysis: A BERT-based approach to named entity recognition (NER) and classification for Malayalam language," *International Journal of Information Technology*, vol. 17, no. 7, pp. 4021-4027, 2025. https://doi.org/10.1007/s41870-024-02248-4

[3] Q. Zhang, Z. Song, D. Wang, Y. Cai, M. Bi, and M. Zuo, "Named entity recognition and coreference resolution using prompt-based generative multimodal," *Complex & Intelligent Systems*, vol. 12, p. 10, 2026. https://doi.org/10.1007/s40747-025-02122-1

[4] A. Lien-Talks, "Evaluating natural language processing and named entity recognition for bioarchaeological data reuse," *Heritage*, vol. 9, no. 1, p. 35, 2026. https://www.mdpi.com/2571-9408/9/1/35

[5] S. Emedem, I. E. Onyenwe, and E. G. Onyedinma, "Development of a WAZOBIA-named entity recognition system," *arXiv preprint arXiv:2505.07884*, 2025. https://doi.org/10.48550/arXiv.2505.07884

[6] B. Paul *et al.*, "Ancholik-ner: A benchmark dataset for bangla regional named entity recognition," *Plos One*, vol. 21, no. 2, p. e0342786, 2026. https://doi.org/10.1371/journal.pone.0342786

[7] B. Yang, C. Zhou, S. Li, and Y. Wang, "A Chinese named entity recognition method for landslide geological disasters based on deep learning," *Engineering Applications of Artificial Intelligence*, vol. 139, p. 109537, 2025. https://doi.org/10.1016/j.engappai.2024.109537

[8] A. Benmachiche, A. Sahia, S. O. Boufaida, K. Rais, M. Derdour, and F. Maazouzi, "Enhancing learning recommendations in mooc search engines through named entity recognition," *Education and Information Technologies*, vol. 30, no. 9, pp. 13041-13071, 2025. https://doi.org/10.1007/s10639-024-13308-4

[9] N. Goyal and N. Singh, "Named entity recognition and relationship extraction for biomedical text: A comprehensive survey, recent advancements, and future research directions," *Neurocomputing*, vol. 618, p. 129171, 2025. https://doi.org/10.1016/j.neucom.2024.129171

[10] S. Saha and S. Mandal, "Enhancing library entity classification through machine learning and named entity recognition," *DESIDOC Journal of Library & Information Technology*, vol. 46, no. 1, pp. 84-93, 2026. https://doi.org/10.14429/djlit.20328

[11] X. Hui, "Research on medical named entity recognition technology based on prompt BioMRC model for deep NLP algorithm," *International Journal of Big Data Intelligent Technology*, vol. 6, no. 1, pp. 118-125, 2025.

[12] P. Sarmah, M. P. Lahkar, S. Kalita, and U. Sharma, "Recent developments in named entity recognition techniques for Indian languages: A comprehensive review," *Natural Language Processing Journal*, p. 100199, 2026. https://doi.org/10.1016/j.nlp.2026.100199

[13] K. Aziz *et al.*, "Advancing Urdu named entity recognition: Deep learning for aspect targeting," *Complex & Intelligent Systems*, vol. 11, p. 489, 2025. https://doi.org/10.1007/s40747-025-02066-6

[14] Z. Zhou, L. Wei, and H. Luan, "Deep learning for named entity recognition in extracting critical information from struck-by accidents in construction," *Automation in Construction*, vol. 173, p. 106106, 2025. https://doi.org/10.1016/j.autcon.2025.106106

[15] M. S. Tiwana, A. Braytee, F. Karimi, M. Bandara, T. Jan, and M. Prasad, "DisasterSense: Event classification in social media using RAKE-based pattern mining and dictionary-driven named entity recognition," *Machine Learning with Applications*, p. 100882, 2026. https://doi.org/10.1016/j.mlwa.2026.100882

[16] Z. Rafiq, M. Wasim, F. S. Shaikh, N. Aldhafferi, and A. Alqahtani, "Urdu-NERD: Urdu named entity recognition with BiGRU-based deep learning architecture," *PeerJ Computer Science*, vol. 12, p. e3678, 2026.

[17] A. Hussein *et al.*, "ECHANT: An efficient context hierarchical attention network for named entity recognition," *International Journal of Data Science and Analytics*, vol. 22, p. 102, 2026. https://doi.org/10.1007/s41060-026-01083-z

[18] B. Liu, Z. Jiang, X. Qian, M. Tao, S. Gao, and X. Li, "Dynamic lexicon integration and update in BERT for enhanced named entity recognition in Chinese legal contracts," *International Journal of Machine Learning and Cybernetics*, vol. 17, no. 3, p. 96, 2026. https://doi.org/10.1007/s13042-025-02961-x

[19] S. Archana and J. Prakash, "Leveraging contextual information in biomedical named entity recognition: Determining the optimal depth of context," *International Journal on Artificial Intelligence Tools*, vol. 35, no. 01, p. 2540012, 2026. https://doi.org/10.1142/S0218213025400123

[20] P. N. Ahmad, M. S. Anwar, S. Masood, A. U. Rehman, and M. Zubair, "Real-time named entity recognition from textual electronic clinical records in cancer therapy using low-latency neural networks," *Big Data*, vol. 14, no. 2, pp. 137-154, 2026. https://doi.org/10.1177/2167647X251409135

[21] B. Gunapriya, A. Thirumalraj, V. Anusuya, B. P. Kavin, and G. H. Seng, *A smart innovative pre-trained model-based QDM for weed detection in soybean fields. In Advanced Intelligence Systems and Innovation in Entrepreneurship*. Hershey, PA, USA: IGI Global Scientific Publishing, 2024.

[22] S. Stephe, V. Revathi, B. Gunapriya, and A. Thirumalraj, *Blockchain-based private AI model with RPOA based sampling method for credit card fraud detection. In Sustainable Development Using Private AI*. Boca Raton, FL, USA: CRC Press, 2024.

[23] Weibo, "Weibo-version-1(Multimodel-Supplement)," 2024. https://www.kaggle.com/datasets/yangzhou32/weibo-version-1multimodel-supplement

[24] Youku, "Youku-mplug: A 10 million large-scale chinese video-language dataset for pre-training and benchmarks," 2023. https://github.com/X-PLUG/Youku-mPLUG

[25] Chinese Literature, "Chinese-Literature-NER-RE-Dataset: A discourse-level named entity recognition and relation extraction dataset for Chinese literature text [Data set]. GitHub," 2017. https://github.com/lancopku/Chinese-Literature-NER-RE-Dataset

Edelweiss Applied Science and Technology
ISSN: 2576-8484
Vol. 10, No. 10: 293-306, 2026
DOI: 10.55214/2576-8484.v10i10.13736
© 2026 by the authors; licensee Learning Gate