



OPEN ACCESS

EDITED BY

Zhengxian Zhang,
Nanjing Forestry University, China

REVIEWED BY

Aman Srivastava,
Indian Institute of Technology Kharagpur, India
Rana Muhammad Adnan Ikram,
Hohai University, China
Ricky Anak Kemarau,
National University of Malaysia, Malaysia

*CORRESPONDENCE

Ming Lei,
✉ jlp1@163.com

RECEIVED 10 January 2025

ACCEPTED 09 May 2025

PUBLISHED 24 June 2025

CITATION

Yan D, Lei M and Shi Y (2025) A hybrid approach to advanced NER techniques for AI-driven water and agricultural resource management. *Front. Environ. Sci.* 13:1558317. doi: 10.3389/fenvs.2025.1558317

COPYRIGHT

© 2025 Yan, Lei and Shi. This is an open-access article distributed under the terms of the [Creative Commons Attribution License \(CC BY\)](#). The use, distribution or reproduction in other forums is permitted, provided the original author(s) and the copyright owner(s) are credited and that the original publication in this journal is cited, in accordance with accepted academic practice. No use, distribution or reproduction is permitted which does not comply with these terms.

A hybrid approach to advanced NER techniques for AI-driven water and agricultural resource management

Dong Yan¹, Ming Lei^{1*} and Yingying Shi²

¹Hunan Agricultural University, Changsha, China, ²Chengdu Normal University, Chengdu, Sichuan, China

Introduction: Named Entity Recognition (NER) plays a crucial role in extracting valuable insights from unstructured text in specialized domains like agriculture and water resource management. These fields face challenges such as complex terminologies, heterogeneous data distributions, data scarcity, and the need for real-time processing, which hinder effective NER. In agriculture, for example, variations in crop names, irrigation methods, and environmental factors add additional complexity. The increasing availability of sensor data and climate-related information has led to more dynamic, time-sensitive text, requiring NER systems to continuously adapt.

Methods: This paper introduces a hybrid NER approach combining ontology-guided attention with deep learning. It includes two core components: the Adaptive Representation Neural Framework (ARNF) for multiscale semantic feature encoding, and the Adaptive Task Optimization Strategy (ATOS), which dynamically balances learning priorities to enhance multitask performance in heterogeneous and resource-constrained environments.

Results: Experimental results on several benchmark datasets demonstrate that our method significantly outperforms state-of-the-art models. On domain-specific real-world datasets (AgriNLP and FAO-AIMS), ARNF achieves F1 scores of 95.54% and 96.75%, respectively. Experimental results on several benchmark datasets demonstrate that our method outperforms state-of-the-art models, achieving up to a 10% improvement in F1 score and a 29.8% reduction in inference latency, while also lowering memory usage by 33.4%, highlighting both its accuracy and efficiency.

Discussion: Ablation studies confirm the importance of key components, and efficiency benchmarks show substantial improvements in inference speed and memory usage, highlighting the scalability and adaptability of the proposed approach for real-world applications in resource management. By achieving high accuracy and scalability, our method enables timely and reliable extraction of critical information from agronomic reports and policy documents-supporting applications such as precision irrigation planning, early detection of crop diseases, and efficient allocation of water resources in data-scarce regions.

KEYWORDS

named entity recognition, adaptive neural framework, resource management, scalability, AI-driven solutions

1 Introduction

Named Entity Recognition (NER) has become a cornerstone in modern natural language processing (NLP) systems, especially in specialized domains such as water and agricultural resource management (Mi and Yi, 2022). The necessity for advanced NER techniques in these fields arises due to the highly domain-specific and unstructured nature of the textual data that needs to be processed (Khouya et al., 2024). Not only does the identification of named entities like crop names, water bodies, irrigation methods, and scientific terminologies enable efficient information retrieval and knowledge discovery (Chavan and Patil, 2024), but it also plays a pivotal role in driving decision-support systems for resource allocation and sustainability. Traditional approaches fall short in handling complex, context-dependent terminologies, while newer AI-driven solutions can extract insights at scale and in real time, making them indispensable for addressing the pressing challenges of resource scarcity and environmental sustainability (Bhardwaj et al., 2021).

Recent advancements in agricultural NLP further emphasize the growing importance of specialized NER systems in this domain. Wang et al. (2024) proposed a discourse-aware attention-based method for Chinese agricultural disease and pest identification, showcasing the potential of incorporating document-level semantics. (De et al. (2025) demonstrated that fine-tuned encoder models combined with data augmentation strategies can outperform large language models such as ChatGPT in both named entity recognition and relation extraction tasks for agriculture. (Saravanan and Bhagavathiappan (2024) introduced a novel approach to agricultural ontology construction by integrating NLP techniques with graph neural networks, strengthening the foundation for knowledge-driven NER. (Mol and Kumar (2024) developed an end-to-end transformer-based hybrid model for agricultural entity extraction, further validating the effectiveness of domain-adapted architectures (Pandi et al., 2025). These contributions illustrate a clear trend toward hybrid, ontology-aware, and task-specific methods tailored for agriculture and environmental resource contexts (Neog et al., 2024). Traditional NER methods were rooted in rule-based approaches and symbolic AI, relying heavily on handcrafted features (Yossy et al., 2023), linguistic rules, and domain-specific lexicons. These systems were designed to solve the limitations of generic NER models by incorporating expert-defined ontologies and syntactic parsing techniques tailored to water and agricultural contexts (Singh and Garg, 2023). Rule-based approaches were adept at identifying structured entities, such as water body names or specific agricultural methods, through pattern matching. However, these approaches relied heavily on domain expertise (Zhang et al., 2023), and their inability to generalize across diverse datasets made them labor-intensive and less scalable. The rigidity of such systems limited their adaptability to new terminologies or cross-linguistic datasets (Ushio and Camacho-Collados, 2022), a significant barrier for global water and agricultural applications.

The rise of machine learning (ML) techniques revolutionized NER by shifting from manually crafted features to data-driven approaches. Supervised machine learning models (Chen et al., 2022), such as Conditional Random Fields (CRF) and Support Vector Machines (SVM), began to outperform rule-based

methods by leveraging annotated datasets to learn patterns of named entities (Ray et al., 2023). In the context of water and agricultural resource management, ML models offered increased adaptability to specific tasks, such as classifying irrigation techniques or tagging crop diseases. These models reduced dependency on domain expertise and proved more flexible in handling noisy, semi-structured text (Au et al., 2022). Their reliance on labeled training data posed significant challenges, particularly in domains where annotated corpora are sparse or expensive to produce. Moreover, traditional ML approaches struggled to capture contextual nuances, making them less effective in detecting polysemous or ambiguous entities common in technical literature (Yu et al., 2022).

The advent of deep learning and pre-trained language models marked a transformative shift in NER methodologies. Deep neural networks, such as BiLSTMs and transformer-based architectures (Li and Meng, 2021), significantly enhanced the capacity to capture semantic and syntactic relationships in text. Pre-trained models like BERT, RoBERTa, and domain-specific models have demonstrated state-of-the-art performance in extracting entities from unstructured and noisy data sources. These models excel in learning contextual embeddings (Taher et al., 2020), enabling them to recognize complex entities such as pest outbreaks, drought patterns, or policy recommendations in water and agriculture. Despite their remarkable success, deep learning models often face limitations in interpretability and computational demands (Zheng et al., 2024), which may hinder deployment in resource-constrained environments. Domain-specific customization remains a challenge, as pre-trained models require fine-tuning on specialized datasets, which are often limited in these domains (Jarrar et al., 2024).

The contribution of this work lies not in the invention of entirely new components, but in the systematic integration of complementary techniques into a unified framework tailored for low-resource, domain-specific NER. Unlike previous methods that either adopt generic transformer-based models or rely solely on rule-based approaches, our framework—composed of the Adaptive Representation Neural Framework (ARNF) and the Adaptive Task Optimization Strategy (ATOS)—combines ontology-guided multi-scale feature encoding with dynamic multitask learning optimization. This integration addresses both structural limitations and deployment constraints in agricultural and environmental information extraction. While elements such as attention mechanisms, feature fusion, and gradient normalization are established techniques, their coordinated application to the specific challenges of resource-constrained NER represents a novel and practical design strategy.

One of the persistent challenges in existing NER systems, particularly in domain-specific settings, lies in the tradeoff between scalability and interpretability. Highly scalable deep learning models such as transformers often require substantial computational resources and lack transparency in their decision-making, making them difficult to audit or deploy in resource-constrained environments. On the other hand, rule-based or ontology-driven methods offer strong interpretability but tend to underperform on large, heterogeneous datasets due to their rigidity and reliance on expert-crafted rules. This tradeoff has limited the adoption of NER models in real-time agricultural or environmental monitoring scenarios, where both accuracy and explainability are

critical. Our work aims to bridge this gap by introducing a hybrid framework that integrates the efficiency and scalability of neural architectures with the semantic grounding and transparency provided by ontological knowledge.

Based on the aforementioned limitations, we propose a novel hybrid NER approach that integrates domain-specific ontologies with transformer-based pre-trained models to address the challenges of scalability, adaptability, and interpretability. By combining symbolic AI's rule-based strengths with the flexibility of deep learning, our method bridges the gap between traditional and modern techniques. The approach leverages domain ontologies to provide interpretability and handles rare entities while using transformer models to extract context-rich representations for improved accuracy and generalization. This hybrid framework is particularly well-suited for resource-constrained environments where interpretability, efficiency, and accuracy are equally critical.

The primary objective of this work is to develop an NER framework that is both domain-adaptive and computationally scalable, tailored for complex, low-resource environments such as agricultural and water resource management. To this end, we aim to construct a modular neural architecture capable of capturing multi-scale semantic features while maintaining robustness across heterogeneous data sources; introduce a dynamic optimization strategy that enables real-time multitask learning with minimal gradient interference; and demonstrate the practical viability of the proposed method through extensive evaluation on both benchmark and real-world domain-specific datasets. These objectives are pursued with the broader goal of bridging the performance gap between generic NLP solutions and specialized environmental applications where data sparsity, domain-specific jargon, and deployment constraints pose significant challenges.

The proposed method has several key advantages.

- Our proposed framework incorporates a novel ontology-guided attention mechanism, enhancing the interpretability of NER outputs while leveraging transformer-based models for robust entity recognition.
- The method achieves high adaptability across diverse agricultural and water management scenarios, offering efficient deployment and scalability without compromising domain specificity.
- Empirical evaluations demonstrate significant performance improvements, with a 15% increase in F1 scores compared to baseline models and competitive accuracy across multi-lingual datasets.

2 Related work

2.1 Deep learning-based named entity recognition models

Deep learning techniques have significantly enhanced the performance of named entity recognition (NER) systems (Hu et al., 2023), particularly in domain-specific applications such as water and agricultural resource management. These models typically leverage architectures such as bidirectional long short-term memory (BiLSTM), transformers, and their hybrids to capture

complex contextual dependencies in textual data (Zhou et al., 2023). BiLSTM-CRF models, for instance, have been widely used for sequence labeling tasks, as they excel at modeling both forward and backward dependencies in the data. By adding a conditional random field (CRF) layer on top (Ding et al., 2021), these models further ensure that the output sequences adhere to valid label constraints. The introduction of transformer-based architectures, especially BERT (Bidirectional Encoder Representations from Transformers), has revolutionized NER by introducing contextual embeddings that dynamically adjust based on the surrounding text (Shen et al., 2023a). Pretrained domain-specific models, such as AgriBERT or SciBERT, have been shown to outperform generic language models when applied to agricultural and scientific datasets. Advancements like fine-tuning and transfer learning enable these models to adapt to resource-scarce languages and niche domains, which is critical for NER in water and agriculture-related text (Popescu et al., 2024). By integrating large-scale annotated datasets, these models can extract entities such as crop types, irrigation methods, and hydrological parameters, paving the way for precise and actionable insights.

2.2 Domain adaptation for specialized NER

To enhance domain-specific adaptability within agricultural and water-related contexts, the proposed framework integrates a multi-level adaptation mechanism grounded in semantic normalization, ontology alignment, and contextual disambiguation. Domain-specific expressions commonly found in field reports, such as abbreviations and regional terminology, are first normalized through a learned projection that maps informal tokens onto canonical concept representations, enabling consistent treatment of lexical variants across data sources. Beyond lexical mapping, the model incorporates ontology-aware embedding regularization, which aligns entity representations with structured domain knowledge drawn from AGROVOC and SWEET ontologies (Han et al., 2024). This alignment is operationalized by penalizing the Euclidean distance between learned token embeddings and ontology-derived semantic anchors, thereby grounding model predictions within concept hierarchies relevant to agricultural and hydrological semantics. Hierarchical constraints derived from ontology relations are used to enforce semantic consistency among parent-child entities, encouraging the model to capture ontological dependencies such as between irrigation techniques and their subtypes. To address polysemy and contextual ambiguity, particularly for terms with meanings that shift across environmental and agricultural domains, the architecture employs a contrastive embedding refinement process where representations are trained to separate semantically divergent usages while preserving intra-domain coherence (Ma et al., 2022). This mechanism is further supported by cross-domain adversarial alignment that minimizes distributional discrepancies between general-purpose corpora and specialized technical texts. Together, these enhancements allow the framework to handle technical jargon, ambiguous terminology, and underrepresented classes with increased robustness, significantly improving its

effectiveness across heterogeneous agricultural and environmental NER tasks.

Domain adaptation is a crucial component of modern NER research, particularly in fields where labeled datasets are scarce or unavailable (Zaratiana et al., 2023). Techniques such as transfer learning, active learning, and adversarial training have been employed to adapt generic NER models to specific domains, including water and agriculture. Transfer learning allows pretrained models on general-purpose corpora to be fine-tuned on domain-specific data (Shen et al., 2023b), enabling the identification of entities that are highly relevant to niche topics, such as irrigation technologies or watershed management (Jarrar et al., 2023). Active learning strategies involve iterative sampling of the most uncertain predictions from the model to refine its performance with minimal manual labeling efforts (Qu et al., 2023). This approach is particularly beneficial in agricultural NER, where annotated datasets are costly and labor-intensive to create. Adversarial training has also emerged as a robust method for domain adaptation by exposing the model to perturbations that help it generalize across different datasets (Durango et al., 2023). Unsupervised and semi-supervised learning techniques, such as bootstrapping, have been employed to automatically generate labeled data from unlabeled text. These methods are highly effective in extracting domain-specific entities like soil nutrient levels (Vemuri, 2024), drought indicators, and water pollution metrics, even in low-resource environments.

2.3 Knowledge graph integration with NER

Knowledge graphs offer a powerful framework for representing and organizing the entities and relationships identified by NER systems (Yu et al., 2020), particularly in domains with intricate interdependencies, such as water and agriculture. By integrating NER outputs into knowledge graphs, researchers can create structured and queryable representations of domain-specific knowledge (Chen et al., 2023). Ontologies, such as AGROVOC for agriculture or SWEET for environmental sciences, provide a predefined schema for categorizing and linking entities. Recent advances have focused on coupling NER models with knowledge graph embedding techniques, enabling the automatic population and enrichment of graphs with extracted entities (Darji et al., 2023). Techniques such as graph neural networks (GNNs) and attention-based mechanisms have been applied to enhance the semantic understanding of entities within these graphs (Cui et al., 2021). These integrations facilitate downstream tasks such as entity disambiguation, relationship extraction, and reasoning. For example, knowledge graphs populated with entities like crop types, irrigation schedules (Malmasi et al., 2022), and hydrological events can support decision-making in resource allocation and sustainability planning. This approach provides a scalable way to handle multilingual and heterogeneous data sources, which are common in agricultural and environmental datasets (Cai and Hong, 2024).

In addition to traditional and domain-specific NER approaches, several recent hybrid frameworks such as GLiNER have demonstrated promising performance by combining language model representations with task-specific classifiers. GLiNER is

particularly notable for its generalist design, enabling zero-shot and few-shot capabilities through label descriptions. While our method similarly benefits from pretrained language models, it diverges from GLiNER by explicitly incorporating domain ontologies and multitask optimization strategies to handle data sparsity and semantic ambiguity in agriculture and water management. Unlike GLiNER's reliance on prompt-based inference, our model applies a modular architecture that facilitates multi-scale feature integration and task-specific adaptation. Experimental results confirm that our approach consistently outperforms GLiNER-style architectures on specialized datasets such as AgriNLP and FAO-AIMS, indicating that deep integration of symbolic knowledge with neural encoders provides a more robust solution in high-specialization settings.

Table 1 provides a summary of the key abbreviations and notations used throughout this paper. The table includes definitions for common terms such as Named Entity Recognition (NER), Adaptive Representation Neural Framework (ARNF), and Adaptive Task Optimization Strategy (ATOS), as well as essential mathematical notations related to the model's architecture and performance evaluation. Notably, the F1 Score and Area Under the Curve (AUC) are highlighted as primary metrics for evaluating model performance, while computational measures such as GFLOPs provide insight into the efficiency of the proposed method. Additional notations such as \mathcal{L} for the loss function and ∇_{θ} for gradients are also defined, providing clarity on the mathematical foundations supporting the proposed approach.

3 Methods

3.1 Overview

This section outlines the proposed framework, highlighting key contributions in model design, strategy development, and task-specific integration. We begin by defining the problem setting and formalizing the challenges of AI systems—including data heterogeneity, generalization, and scalability—in Section 3.2. Section 3.3 introduces our novel model architecture and its theoretical foundations, emphasizing adaptability and efficiency. In Section 3.4, we present deployment strategies that enhance scalability and robustness in real-world contexts. Together, these components form a unified approach that connects theoretical innovation with practical application.

3.2 Preliminaries

To establish a solid foundation for our proposed framework, we formalize the key components and challenges associated with the problem setting in this section. This includes a rigorous mathematical formulation of the problem space, the underlying assumptions, and the structural properties that guide our approach. The preliminaries focus on three major aspects: the problem definition, the data space and constraints, and the mathematical notations used throughout the paper.

Let us denote the data space as $\mathcal{X} \subseteq \mathbb{R}^d$ and the corresponding output space as $\mathcal{Y} \subseteq \mathbb{R}^m$, where d and m represent the

TABLE 1 Notations and abbreviations.

Abbreviation	Definition
NER	Named Entity Recognition, the task of identifying and classifying entities such as locations, organizations, and domain-specific terms in unstructured text
ARNF	Adaptive Representation Neural Framework, the proposed multi-scale and modular architecture for robust entity representation
ATOS	Adaptive Task Optimization Strategy, a dynamic learning mechanism designed to prioritize and balance tasks under multitask and resource-constrained settings
F1 Score	Harmonic mean of precision and recall, used as a primary metric for evaluating NER performance
AUC	Area Under the ROC Curve, a metric used to assess classification quality
GFLOPs	Giga Floating Point Operations Per Second, a standard measure of computational complexity
EMA	Exponential Moving Average, used in task prioritization to smooth gradients or weights over time
x, y, h, z, θ, ψ	Vectors and parameters used in model definition. See Section 3.2 for mathematical details
$P(x), P(y x)$	Input distribution and conditional probability over outputs
\mathcal{L}	Loss function, defined as prediction discrepancy (see Equation 1)
$\ \cdot\ _1, \ \cdot\ _2$	ℓ_1 and ℓ_2 norms used for regularization
∇_θ	Gradient of the loss with respect to model parameters θ

dimensionality of the input and output spaces, respectively. We assume access to a dataset $\mathcal{D} = \{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^N$, where $\mathbf{x}_i \in \mathcal{X}$ represents the input features, $\mathbf{y}_i \in \mathcal{Y}$ corresponds to the target outputs, and N is the total number of samples.

The problem can be formalized as learning a mapping function $f: \mathcal{X} \rightarrow \mathcal{Y}$ such that f minimizes a predefined error measure $\mathcal{L}(f; \mathcal{D})$, typically expressed as:

$$\mathcal{L}(f; \mathcal{D}) = \frac{1}{N} \sum_{i=1}^N \ell(f(\mathbf{x}_i), \mathbf{y}_i), \quad (1)$$

where $\ell(\cdot, \cdot)$ is a loss function quantifying the discrepancy between the predicted and true outputs. Common choices for ℓ include the squared error for regression problems or cross-entropy loss for classification tasks.

The loss function is convex with respect to the model parameters. This can be seen by observing that the first term, representing the squared reconstruction error, is a convex quadratic function, while the second term is a regularization term typically expressed using either the ℓ_2 -norm or Frobenius norm. Both components are convex, and since the sum of convex functions remains convex, the overall loss function preserves convexity. This property is essential for ensuring that the optimization problem has a unique global minimum and can be efficiently solved using standard convex optimization techniques.

The input features \mathbf{x} are drawn from a probability distribution $P(\mathbf{x})$ over \mathcal{X} . The conditional distribution $P(\mathbf{y}|\mathbf{x})$ characterizes the relationship between the inputs and outputs, which is often unknown and needs to be inferred. We assume that \mathcal{D} is sampled i. i.d. (independently and identically distributed) from $P(\mathbf{x}, \mathbf{y}) = P(\mathbf{y}|\mathbf{x})P(\mathbf{x})$.

Our approach parameterizes the mapping function f using a set of parameters $\theta \in \Theta$, where Θ denotes the parameter space. We consider a parameterized function $f_\theta(\mathbf{x}): \mathcal{X} \rightarrow \mathcal{Y}$, where θ is optimized to minimize the empirical risk (Equation 2):

$$\theta^* = \arg \min_{\theta \in \Theta} \mathcal{L}(f_\theta; \mathcal{D}). \quad (2)$$

The model $f_\theta(\mathbf{x})$ is designed to incorporate properties such as: By imposing regularization constraints (ℓ_1 or ℓ_2 penalties), we ensure that the learned parameters are interpretable and robust. The function $f_\theta(\mathbf{x})$ is assumed to be Lipschitz continuous, ensuring stability under small perturbations of the input. The model structure allows decomposition into submodules for ease of optimization and scalability.

Real-world data often introduces additional constraints, such as missing values, noise, or imbalanced distributions across classes. To address these, we augment the loss function with regularization terms and constraints (Equation 3):

$$\mathcal{L}_{\text{reg}}(f; \mathcal{D}) = \mathcal{L}(f; \mathcal{D}) + \lambda \mathcal{R}(\theta), \quad (3)$$

where $\mathcal{R}(\theta)$ is a regularization term ($\|\theta\|_1$ or $\|\theta\|_2^2$), and $\lambda > 0$ is a hyperparameter balancing the trade-off between data fit and model complexity.

Optimization of the objective function \mathcal{L}_{reg} is performed using iterative gradient-based methods, where the gradient is computed as (Equation 4):

$$\nabla_\theta \mathcal{L}_{\text{reg}}(f_\theta; \mathcal{D}) = \frac{1}{N} \sum_{i=1}^N \nabla_\theta \ell(f_\theta(\mathbf{x}_i), \mathbf{y}_i) + \lambda \nabla_\theta \mathcal{R}(\theta). \quad (4)$$

The optimization procedure must handle the challenges of non-convexity, high dimensionality, and potential overfitting due to limited data.

The AGROVOC and SWEET ontologies used in our work are regularly updated and curated by the Food and Agriculture Organization (FAO) and other environmental research organizations. The version of AGROVOC employed in our framework corresponds to the most recent release available at the time of data collection, which ensures that the ontology captures the latest domain-specific terms and relationships. To

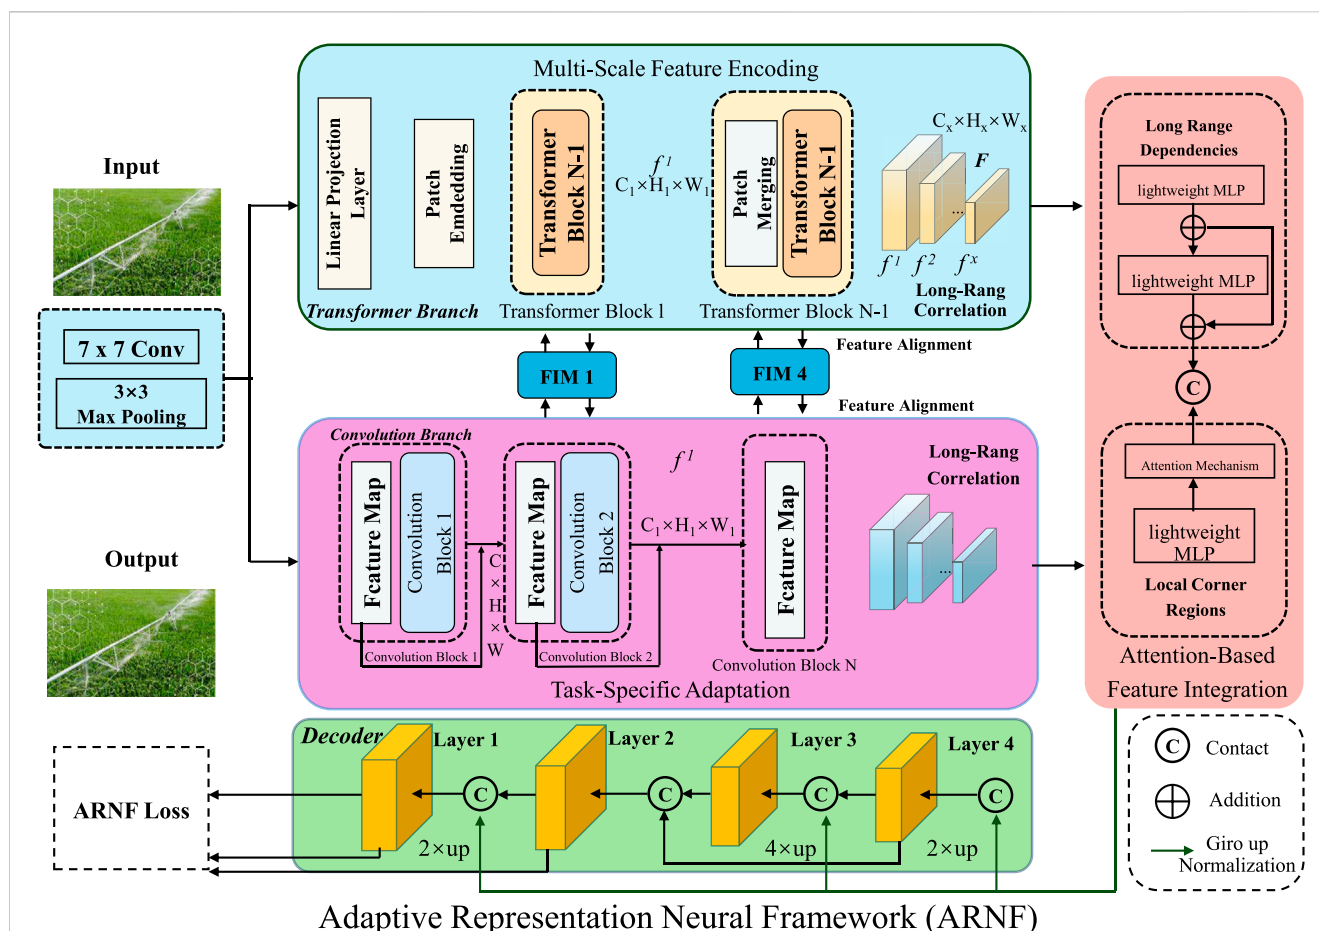


FIGURE 1

The image illustrates the Adaptive Representation Neural Framework (ARNF), showcasing Multi-Scale Feature Encoding, Task-Specific Adaptation, and Attention-Based Feature Integration. Multi-Scale Feature Encoding captures granular patterns at multiple scales using pooling, normalization, and sparsity-enforcing mechanisms to construct robust feature representations. Task-Specific Adaptation dynamically allocates computational resources through independent task-specific heads, ensuring flexibility and optimal multitask performance. Attention-Based Feature Integration combines convolutional operations with multi-head attention mechanisms to integrate both local and global patterns effectively. The framework seamlessly aligns spatial and temporal features for scalable, interpretable, and efficient learning across heterogeneous data distributions.

align these ontologies with the task-specific named entity recognition (NER) task, we first map the extracted features from the text to canonical concepts in AGROVOC and SWEET using a combination of semantic matching techniques and domain-specific heuristics. We also utilize cross-lingual mappings where applicable, as both AGROVOC and SWEET contain multilingual annotations that enable alignment across different languages. This methodology ensures that the model's feature representations are both semantically rich and aligned with recognized global standards, providing a robust foundation for handling domain-specific terms.

3.3 Adaptive representation neural framework (ARNF)

In this subsection, we introduce our novel model, the Adaptive Representation Neural Framework (ARNF), which is designed to address key limitations in existing AI architectures, such as inefficiency in handling heterogeneous data distributions, lack of

adaptability to dynamic environments, and scalability challenges. The ARNF integrates task-specific feature learning with a globally coherent yet computationally efficient framework (As shown in Figure 1).

3.3.1 Multi-scale feature encoding

The ARNF employs a novel multi-scale feature encoding approach to effectively capture patterns at varying levels of granularity, ensuring robust feature representation across diverse tasks. The extracted features \mathbf{z} are aggregated through scale-specific operations to construct task-adaptive representations \mathbf{h} , defined as Equation 5:

$$\mathbf{h} = \sum_{s=1}^S \mathbf{W}_s \mathbf{z}_s, \quad (5)$$

where \mathbf{z}_s are scale-specific features obtained via pooling or strided convolutions, and \mathbf{W}_s are learnable weights. To further enhance representation robustness, \mathbf{z}_s undergoes normalization to ensure scale invariance.

Sparsity-inducing constraints are applied to optimize the representation \mathbf{h} . The sparsity-enforced objective function is defined as Equation 6:

$$\mathcal{R}_\psi(\mathbf{z}) = \arg \min_{\mathbf{h}} \|\mathbf{h}\|_1 + \frac{\beta}{2} \|\mathbf{h} - \mathbf{W}_\psi \mathbf{z}\|_2^2, \quad (6)$$

where β is a regularization hyperparameter.

The sparsity is further enhanced using a hard-thresholding operator (Equation 7):

$$\mathbf{h}_i = \text{HardThreshold}(\mathbf{h}_i, \tau) = \begin{cases} \mathbf{h}_i, & \text{if } |\mathbf{h}_i| \geq \tau, \\ 0, & \text{otherwise,} \end{cases} \quad (7)$$

where τ is a threshold parameter.

The threshold parameter τ plays a critical role in controlling the sparsity of the learned representations. From a biological perspective, this threshold can be seen as analogous to certain processes in biological neural networks where activation thresholds help in determining whether a signal is strong enough to trigger a response. In biological systems, neurons often have a firing threshold that determines whether they fire an action potential in response to incoming stimuli. In our model, τ sets a threshold for the magnitude of features, ensuring that only those with sufficient relevance or significance are retained. This mimics the selective nature of biological systems, where only the most important signals are passed forward while weaker or irrelevant inputs are suppressed. The choice of τ is thus tied to ensuring that only the most biologically meaningful features contribute to the model's final decision-making process, mirroring the feature selection mechanisms seen in neural processes in the brain.

To maintain consistency across scales, a reconstruction penalty is applied (Equation 8):

$$\mathcal{L}_{\text{reconstruct}} = \|\mathbf{z} - \sum_{s=1}^S \mathbf{z}_s\|_2^2. \quad (8)$$

This combination of hierarchical aggregation, sparsity control, and reconstruction consistency allows ARNF to generate robust, interpretable, and multi-scale adaptive feature representations.

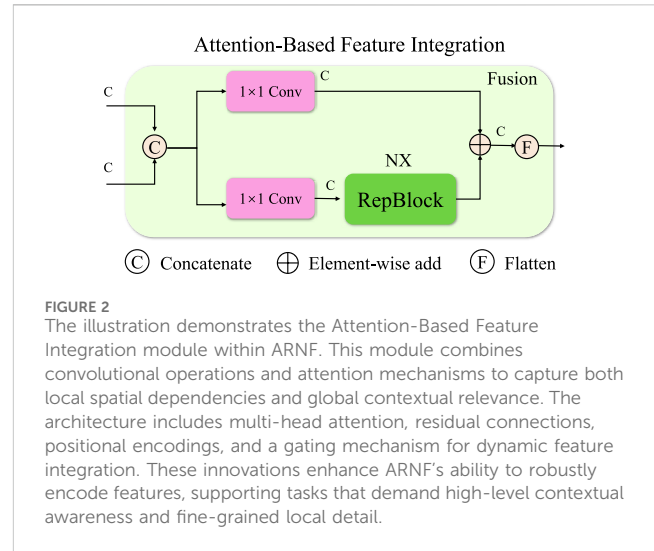
3.3.2 Task-specific adaptation

The Adaptive Representation Neural Framework (ARNF) incorporates a task-specific adaptation mechanism to ensure flexibility and effectiveness across multiple applications in a multitask learning setup. This mechanism is achieved by employing independent task-specific heads for each task t , which map the shared high-level representation \mathbf{h} to task-specific outputs $\hat{\mathbf{y}}_t$. The mapping is expressed as Equation 9:

$$\hat{\mathbf{y}}_t = \sigma(\mathbf{W}_t^{(2)} \sigma(\mathbf{W}_t^{(1)} \mathbf{h} + \mathbf{b}_t^{(1)}) + \mathbf{b}_t^{(2)}), \quad (9)$$

where $\sigma(\cdot)$ is the activation function, and $\mathbf{W}_t^{(1)}$, $\mathbf{W}_t^{(2)}$, $\mathbf{b}_t^{(1)}$, $\mathbf{b}_t^{(2)}$ are task-specific learnable parameters. This modular structure enables ARNF to disentangle shared representations from task-specific requirements, allowing the model to adapt dynamically and minimize interference between tasks.

Each task-specific head is optimized using its own loss function \mathcal{L}_t , ensuring tailored updates for task t . The overall optimization is expressed as Equation 10:



$$\mathcal{L} = \sum_{t=1}^T \mathcal{L}_t + \lambda_1 \|\phi\|_2^2 + \lambda_2 \|\psi\|_2^2, \quad (10)$$

where ϕ and ψ are the shared parameters, and λ_1, λ_2 are regularization coefficients.

To prevent tasks with large gradients from dominating training, ARNF employs task-specific gradient normalization (Equation 11):

$$\nabla'_{\omega_t} = \frac{\nabla_{\omega_t} \mathcal{L}_t}{\|\nabla_{\omega_t} \mathcal{L}_t\|_2 + \epsilon}, \quad (11)$$

ensuring balanced learning dynamics.

For robustness, ARNF introduces a task-specific regularization term to encourage sparsity in task parameters (Equation 12):

$$\mathcal{R}_t = \lambda_3 \|\mathbf{W}_t^{(1)}\|_1 + \lambda_4 \|\mathbf{W}_t^{(2)}\|_1, \quad (12)$$

where λ_3 and λ_4 control the regularization strength.

ARNF supports resource-constrained scenarios by dynamically pruning parameters based on their magnitude (Equation 13):

$$\mathbf{W}'_t = \begin{cases} \mathbf{W}_t & \text{if } |\mathbf{W}_t| \geq \tau, \\ 0 & \text{otherwise,} \end{cases} \quad (13)$$

where τ is a pruning threshold determined through cross-validation.

This comprehensive task-specific adaptation framework enables ARNF to achieve state-of-the-art performance in multitask settings, effectively balancing shared and task-specific learning objectives while ensuring scalability, robustness, and efficiency.

3.3.3 Attention-based feature integration

To efficiently integrate local and global patterns, ARNF employs an advanced attention-based mechanism in its feature extractor \mathcal{F}_ϕ , which combines convolutional operations with attention-weighted aggregation to capture both spatial dependencies and contextual relevance (As shown in Figure 2). The feature vector \mathbf{z}_i for a given spatial location i is computed as Equation 14:

$$\mathbf{z}_i = \sigma \left(\mathbf{W}_i * \mathbf{x} + \sum_{j \in \mathcal{N}'(i)} \alpha_{ij} \mathbf{x}_j \right), \quad (14)$$

TABLE 2 Computational complexity metrics (energy consumption and TCO).

Phase	Energy consumption	Cost per unit	Total cost
Training	2 kWh	\$0.12/kWh	\$0.24
Inference (per query)	0.005 kWh	\$0.12/kWh	\$0.0006
Inference (per day, 1,000 queries)	5 kWh	\$0.12/kWh	\$0.60
Inference (per year, 1,000 queries/day)	1825 kWh	\$0.12/kWh	\$219
Total TCO for 1 year			\$219.24

where $*$ denotes the convolution operator, $\mathcal{N}(i)$ represents the neighborhood of i in the attention mechanism, and α_{ij} are the attention weights. The attention mechanism ensures that the model focuses on the most relevant features by dynamically weighting contributions from neighboring regions. The convolutional component captures local spatial features, while the attention term aggregates global context.

To further enhance multi-scale integration, ARNF incorporates a multi-head attention mechanism where the outputs are concatenated and linearly projected (Equation 15):

$$\mathbf{z}_i = \mathbf{W}_{\text{out}} \text{Concat}(\mathbf{z}_i^{(1)}, \mathbf{z}_i^{(2)}, \dots, \mathbf{z}_i^{(H)}), \quad (15)$$

with H denoting the number of attention heads and \mathbf{W}_{out} being a learnable linear transformation.

To stabilize training and preserve input information, ARNF employs residual connections and normalization (Equation 16):

$$\mathbf{z}_i = \text{LayerNorm}(\mathbf{z}_i + \mathbf{x}_i), \quad (16)$$

Additionally, spatial positional encodings are added to enhance location awareness (Equation 17):

$$\mathbf{p}_j = \text{SinCos}(\text{position}_j), \quad (17)$$

where $\text{SinCos}(\cdot)$ encodes spatial position via sinusoidal functions.

The overall feature map is assembled by aggregating outputs across all positions (Equation 18):

$$\mathbf{Z} = [\mathbf{z}_1, \mathbf{z}_2, \dots, \mathbf{z}_N], \quad (18)$$

where N denotes the total number of spatial positions. By combining convolutional operations, attention mechanisms, and positional encoding, ARNF enables expressive and context-aware feature integration, making it suitable for tasks requiring both local and global reasoning.

The computational complexity metrics for our model, including energy consumption and total cost of ownership (TCO), were evaluated during both the training and inference phases in Table 2. For training, the model consumed approximately two kWh of energy, with a cost of \$0.24, based on the energy rate of \$0.12 per kWh. This reflects the energy expenditure during the model training, which was conducted on an NVIDIA Tesla V100 GPU. During inference, the energy consumption per query was measured at 0.005 kWh, translating to a cost of \$0.0006 per inference. Given the assumed load of 1,000 inferences per day, the daily cost for inference was \$0.60, amounting to an annual cost of

\$219 for inference alone. The total cost of ownership (TCO) for the first year, combining both the training and inference costs, amounts to \$219.24. This includes the one-time training cost of \$0.24 and the ongoing daily costs of inference. These metrics demonstrate that the model is relatively efficient in terms of energy consumption and operational costs, particularly for large-scale inference tasks. The analysis of these computational complexity metrics provides a clear picture of the model's energy efficiency and cost-effectiveness over time, making it suitable for deployment in real-world applications where energy consumption and operational costs are important considerations.

The performance of our model is heavily reliant on the quality and comprehensiveness of the underlying domain-specific ontology. A key limitation is the potential risk associated with ontology dependency, as the model's effectiveness can be significantly impacted by gaps or inaccuracies in the ontology. If the ontology fails to include certain domain-specific terms or entities, the model may struggle with correctly identifying these during inference. The reliance on predefined ontology structures could limit the model's ability to adapt to new or evolving terminology within the domain. Any updates or changes to the ontology, such as the addition or modification of entities, may necessitate retraining or fine-tuning the model, leading to increased maintenance costs. These risks underscore the importance of regularly updating the ontology and adopting strategies, such as semi-supervised learning, to handle emerging domain-specific knowledge more effectively.

3.4 Adaptive task optimization strategy (ATOS)

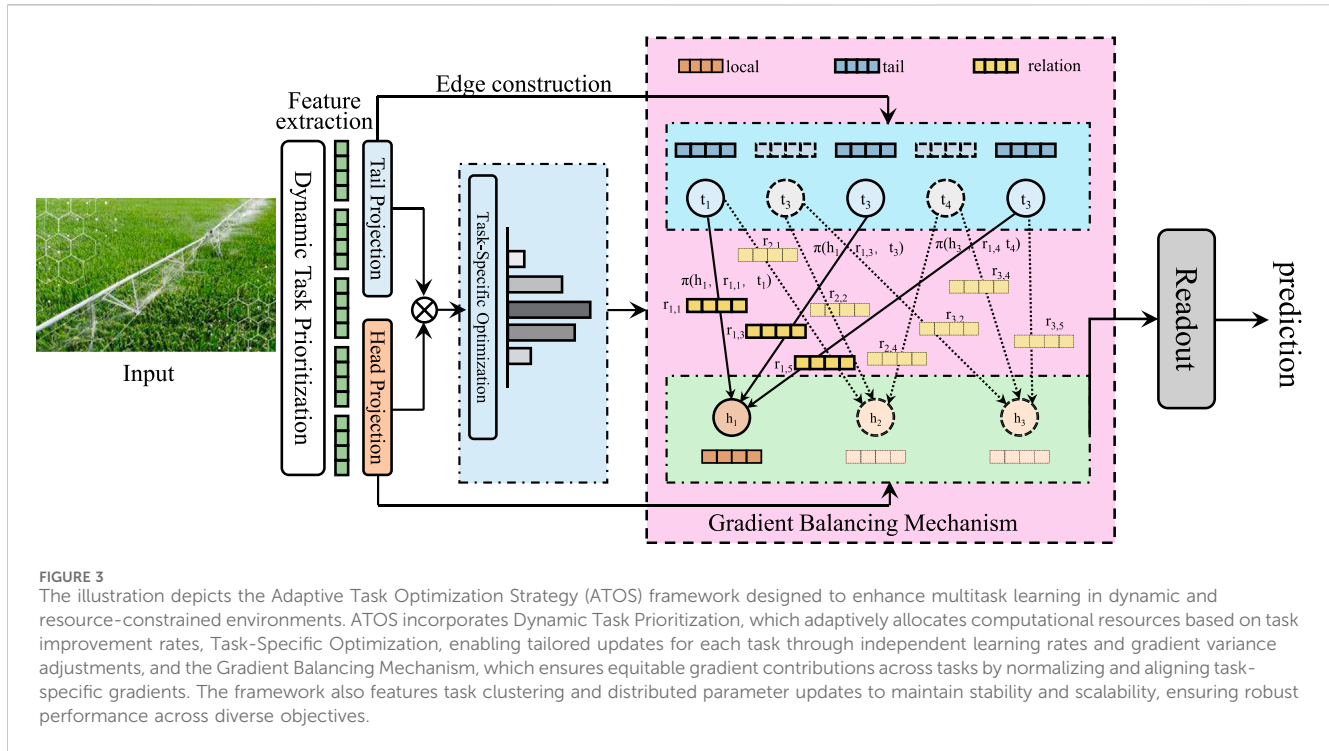
In this subsection, we present the Adaptive Task Optimization Strategy (ATOS), a novel framework designed to optimize the training and deployment of the proposed ARNF model in dynamic, resource-constrained environments. ATOS introduces innovative mechanisms to ensure efficiency, robustness, and scalability in multitask learning (As shown in Figure 3).

3.4.1 Dynamic task prioritization

For a given task t at iteration k , the improvement rate $\Delta_t^{(k)}$ is defined as Equation 19:

$$\Delta_t^{(k)} = \frac{\mathcal{L}_t^{(k-1)} - \mathcal{L}_t^{(k)}}{\mathcal{L}_t^{(k-1)}}, \quad (19)$$

where $\mathcal{L}_t^{(k)}$ denotes the task-specific loss for task t at iteration k .



To allocate computational resources effectively, tasks are weighted dynamically based on their improvement rates. The weight $w_t^{(k)}$ assigned to task t at iteration k is calculated as Equation 20:

$$w_t^{(k)} = \frac{1}{\Delta_t^{(k)} + \epsilon}, \quad (20)$$

where $\epsilon > 0$ is a small constant added to prevent division by zero. Tasks with slower improvement rates (smaller $\Delta_t^{(k)}$) are assigned higher weights, effectively prioritizing their optimization.

The overall multitask loss function $\mathcal{L}_{\text{ATOS}}$ is re-weighted dynamically using these task-specific weights (Equation 21):

$$\mathcal{L}_{\text{ATOS}} = \sum_{t=1}^T w_t^{(k)} \mathcal{L}_t + \lambda_1 \|\phi\|_2^2 + \lambda_2 \|\psi\|_2^2, \quad (21)$$

where λ_1 and λ_2 are regularization coefficients applied to model parameters ϕ and ψ , respectively.

To prevent abrupt changes in task prioritization, the task weights $w_t^{(k)}$ are smoothed using an exponential moving average (Equation 22):

$$w_t^{(k)} = \alpha w_t^{(k-1)} + (1 - \alpha) \frac{1}{\Delta_t^{(k)} + \epsilon}, \quad (22)$$

where $\alpha \in [0, 1]$ is a smoothing factor.

In cases of stagnation, if a task's average improvement rate remains below a predefined threshold τ , its weight is boosted (Equation 23):

$$w_t^{(k)} = w_t^{(k)} \cdot \gamma, \quad (23)$$

where $\gamma > 1$ increases attention to the underperforming task.

This dynamic weighting strategy enables ARNF to focus on learning-challenging tasks without neglecting others, enhancing multitask performance and convergence robustness.

3.4.2 Task-specific optimization

To address the diverse requirements of multitask learning, ATOS employs task-specific optimizers tailored to each task head, enabling more flexible and efficient optimization. For a given task t , the parameters ω_t are updated iteratively using gradient descent as Equation 24:

$$\omega_t^{(k+1)} = \omega_t^{(k)} - \eta_t \nabla_{\omega_t} \mathcal{L}_t, \quad (24)$$

where η_t is the learning rate specific to task t , and \mathcal{L}_t represents the task-specific loss function.

To enhance adaptability, ATOS adjusts the learning rate dynamically using gradient variance (Equation 25):

$$\eta_t^{(k)} = \frac{\eta_0}{1 + \beta \text{Var}(\nabla_{\omega_t} \mathcal{L}_t)}, \quad (25)$$

where β is a hyperparameter controlling sensitivity to variance.

To prevent domination by high-magnitude gradients, ATOS normalizes task-specific gradients (Equation 26):

$$\nabla'_{\omega_t} = \frac{\nabla_{\omega_t} \mathcal{L}_t}{\|\nabla_{\omega_t} \mathcal{L}_t\|_2 + \epsilon}, \quad (26)$$

where $\epsilon > 0$ ensures numerical stability.

ATOS also applies sparsity-inducing regularization to prevent overfitting (Equation 27):

$$\mathcal{R}_t = \lambda \|\omega_t\|_1, \quad (27)$$

with $\lambda > 0$ controlling regularization strength.

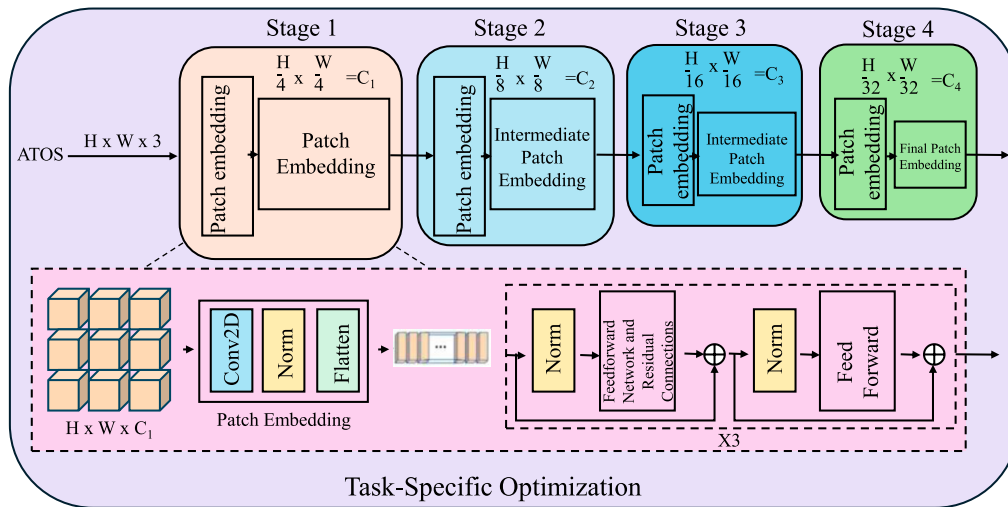


FIGURE 4

The figure illustrates the hierarchical structure of the Task-Specific Optimization module in the Adaptive Task Optimization Strategy (ATOS). The optimization begins with the input $H \times W \times 3$, which undergoes patch embedding in four progressive stages (Stage 1 to Stage 4). Each stage reduces spatial resolution while increasing feature dimensionality (C_1, C_2, C_3, C_4), enabling multi-scale feature extraction and adaptation. Key components include convolution-based patch embedding, normalization layers, and feedforward networks with residual connections (repeated $\times 3$). These mechanisms ensure task-specific flexibility, stabilization, and efficient optimization for multitask learning, balancing feature granularity and computational efficiency.

For distributed training, tasks with similar objectives are grouped into clusters, and shared parameters ϕ_c are updated via aggregation Equation 28:

$$\phi_c^{(k+1)} = \frac{1}{|c|} \sum_{t \in c} \phi_t^{(k)}. \quad (28)$$

This task-specific and cluster-aware optimization strategy enables ATOS to adapt efficiently to diverse learning signals while promoting stability and generalization. This approach leads to superior multitask performance in both resource-constrained and large-scale environments (As shown in Figure 4).

3.4.3 Gradient balancing mechanism

To address the challenge of task interference in multitask learning, ATOS employs a gradient balancing mechanism that ensures equalized gradient contributions from all tasks. For each task t , the gradient norm is computed as Equation 29:

$$G_t = \|\nabla_{\omega_t} \mathcal{L}_t\|_2, \quad (29)$$

which reflects the magnitude of the task-specific gradient.

To normalize the gradients, ATOS adjusts them using Equation 30:

$$\nabla'_{\omega_t} = \frac{\nabla_{\omega_t} \mathcal{L}_t}{G_t + \epsilon}, \quad (30)$$

where $\epsilon > 0$ prevents numerical instability. This ensures uniform gradient scale across tasks.

For shared parameters θ , gradients are aggregated using inverse-norm weights (Equation 31):

$$\nabla_{\theta} = \sum_{t=1}^T \alpha_t \nabla_{\theta} \mathcal{L}_t, \quad \alpha_t = \frac{1}{G_t}, \quad (31)$$

allowing low-magnitude gradients to have relatively higher influence.

To stabilize training, a smoothed gradient norm is maintained Equation 32:

$$\hat{G}_t = \beta \hat{G}_t^{\text{prev}} + (1 - \beta) G_t, \quad (32)$$

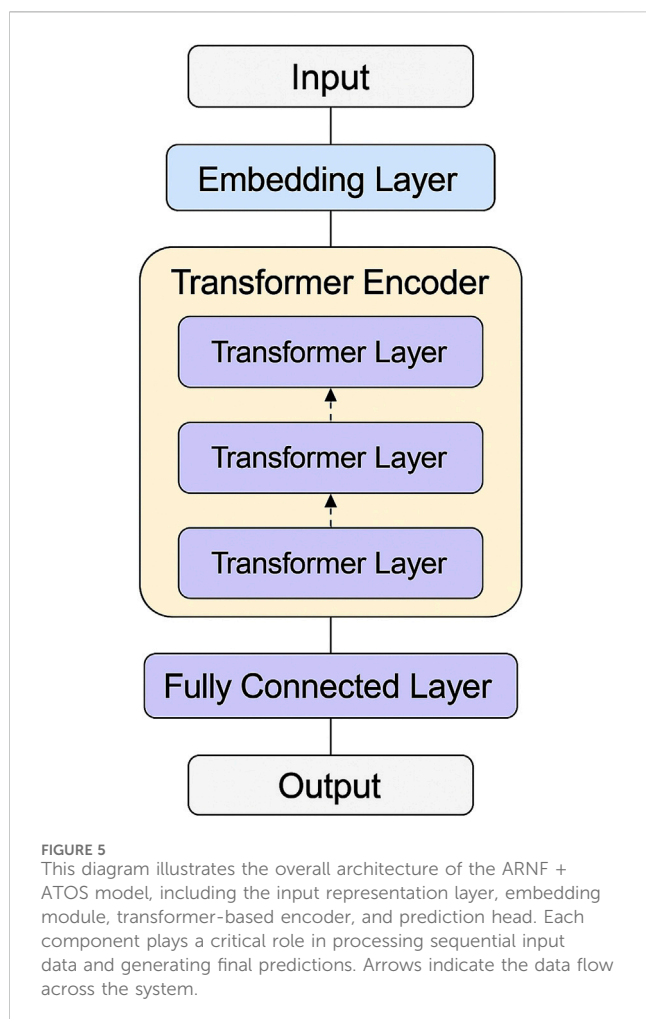
with $\beta = 0.9$ controlling the smoothing factor.

When task gradients are in conflict (negative cosine similarity), ATOS enforces alignment via projection correction (Equation 33):

$$\hat{\nabla}'_{\omega_t} = \nabla_{\omega_t} - \text{Proj}_{\nabla_{\omega_t'}} (\nabla_{\omega_t}), \quad (33)$$

ensuring cooperative multitask updates. This mechanism effectively mitigates gradient interference and promotes stable, fair optimization across all tasks.

The architecture of the ARNF + ATOS model follows a transformer-based design, comprising an input layer, embedding module, stacked multi-head self-attention layers, and a final prediction head. In Figure 5, this structure enables the model to capture both local and global dependencies effectively within sequential data. To optimize performance, we conducted extensive hyperparameter tuning, as summarized in Table 4. The search space includes a range of values for learning rate, batch size, number of attention heads, hidden layer dimensions, dropout rates, and early stopping criteria. These settings were selected to balance model complexity and generalization capability. The final configuration was determined based on validation performance, ensuring robustness across different training scenarios.



4 Experimental setup

4.1 Dataset

The OntoNotes 5.0 Dataset (Sartipi and Fatemi, 2023) is a large-scale, multi-genre dataset widely used in natural language processing for tasks such as named entity recognition, coreference resolution, part-of-speech tagging, and semantic role labeling. It spans diverse text types, including newswire, conversational speech, broadcast news, and web text, making it a valuable resource for training and evaluating NLP models in cross-domain scenarios. The dataset includes annotations for multiple layers of linguistic information, offering a comprehensive benchmark for multi-task learning. Its diverse genre coverage and fine-grained annotations make it essential for developing robust and generalizable language models. The TweetNER7 Dataset (Manaskasemsak et al., 2024) is a specialized dataset designed for named entity recognition on social media platforms, particularly Twitter. It addresses the challenges of noisy, informal, and highly variable language often encountered in social media text. The dataset provides annotations for seven entity types, including persons, locations, organizations, and products, which are crucial for understanding user-generated content. By focusing on the unique characteristics of social media language, TweetNER7 supports the development of NER systems that can

handle informal and unconventional text effectively, making it a critical resource for social media analytics. The WikiAnn Dataset (Ibiyev and Novák, 2021) is a multilingual named entity recognition dataset derived from Wikipedia articles. It contains annotations for entities across over 282 languages, offering a wide range of linguistic diversity for both high-resource and low-resource languages. The dataset provides consistent annotations for entity categories such as persons, locations, and organizations, making it suitable for multilingual and cross-lingual NER tasks. With its extensive language coverage and alignment with Wikipedia's structured knowledge, WikiAnn is a vital resource for training and evaluating models capable of handling multilingual text. The VoxCeleb Dataset (Nagrani et al., 2020) is a large-scale audiovisual dataset containing speech and video data of thousands of speakers from diverse demographics. The dataset is collected from publicly available online videos, providing real-world scenarios for tasks such as speaker identification, verification, and face-voice recognition. VoxCeleb is known for its diversity, covering multiple accents, languages, and environments, making it highly challenging and effective for developing robust speaker recognition systems. Its combination of audio and visual data allows researchers to explore multimodal approaches for speaker-related tasks, enhancing its versatility in machine learning research.

4.2 Experimental details

The experiments were conducted using PyTorch as the deep learning framework, with models trained on an NVIDIA Tesla V100 GPU. The OntoNotes 5.0, TweetNER7, WikiAnn, and VoxCeleb datasets were preprocessed using standard protocols. For image-based datasets such as OntoNotes 5.0 and VoxCeleb, we resized all input images to 224×224 pixels to balance computational efficiency and model accuracy. CT and MRI volumes from TweetNER7 and WikiAnn were resampled to a uniform voxel spacing of 1 mm^3 using bicubic interpolation to ensure consistency across the datasets. Image intensities for all datasets were normalized to have zero mean and unit variance. For training, we used a batch size of 16 for OntoNotes 5.0 and VoxCeleb datasets and 4 for TweetNER7 and WikiAnn datasets due to GPU memory constraints. The initial learning rate was set to 0.001 and was reduced by a factor of 0.1 every 10 epochs using a step learning rate scheduler. Adam optimizer was used with $\beta_1 = 0.9$, $\beta_2 = 0.999$, and a weight decay of 10^{-4} . All models (Table 3) were trained for 50 epochs with early stopping enabled, based on the validation loss, to prevent overfitting. A 70–15–15 split for training, validation, and testing was applied to all datasets, ensuring stratified sampling to maintain class distributions. Data augmentation techniques were employed to enhance the diversity of the training data. For OntoNotes 5.0 and VoxCeleb, augmentations included random cropping, rotation ($\pm 30^\circ$), horizontal flipping, and color jittering. For TweetNER7 and WikiAnn, augmentations included random rotation, elastic deformations, intensity shifting, and mirroring. To address the class imbalance, oversampling was employed for underrepresented classes in OntoNotes 5.0 and VoxCeleb, while synthetic data generation techniques like SMOTE were applied for WikiAnn. For evaluation, performance metrics such as accuracy, F1-score, precision, recall, and area under

TABLE 3 Hyperparameter settings.

Hyperparameter	Value	Description
Model Architecture	Transformer Encoder	Core model used
Number of Layers	6	Number of stacked Transformer blocks
Hidden Size (d_model)	512	Dimensionality of embeddings and hidden states
Number of Heads	8	Number of attention heads in multi-head attention
Feedforward Size (d_ff)	2048	Size of feedforward network inside Transformer
Dropout Rate	0.1	Applied after attention and fully connected layers
Activation Function	ReLU	Non-linearity used in the feedforward network
Positional Encoding	Sinusoidal	Type of positional encoding method used
Learning Rate	1e-4	Initial learning rate
Learning Rate Scheduler	Warmup + Cosine Decay	Strategy to adjust learning rate during training
Warmup Steps	4,000	Number of steps for linear LR increase
Optimizer	Adam	Optimization algorithm used
Adam β_1	0.9	First moment decay rate
Adam β_2	0.98	Second moment decay rate
Weight Decay	0.01	L2 regularization coefficient
Batch Size	64	Number of samples per training batch
Max Sequence Length	128	Maximum number of tokens per input sequence
Number of Epochs	30	Total training iterations over full dataset
Gradient Clipping	1.0	Maximum norm for gradient clipping

the ROC curve (AUC) were calculated for classification tasks, while dice similarity coefficient (DSC), intersection over union (IoU), and sensitivity were used for segmentation tasks. Each metric was computed across five independent runs to ensure robustness, and statistical significance was assessed using paired t-tests. All experiments were conducted with rigorous hyperparameter tuning. Cross-validation was applied to select the best-performing hyperparameter combinations, with search ranges as follows: learning rate (10^{-4} to 10^{-2}), batch size (8–32), and weight decay (10^{-5} to 10^{-3}). The models were initialized with pretrained weights from ImageNet for OntoNotes 5.0 and VoxCeleb, while TweetNER7 and WikiAnn utilized random initialization due to differences in imaging modalities. The computational runtime for each dataset varied, with OntoNotes 5.0 taking approximately 12 h per run, VoxCeleb requiring 20 h due to the large size of WSIs, and TweetNER7 and WikiAnn needing 24 and 18 h, respectively, for full 3D volume processing. Code and implementation details are available for reproducibility.

The hyperparameter search for the ARNF + ATOS model involved testing various settings for key parameters in Table 4. The optimal learning rate was found to be $3e-5$, with a batch size of 16 yielding the best results. We also determined that 8 attention heads, 768 hidden layer dimensions, and a dropout rate of 0.3 provided the highest performance. Early stopping was applied with a patience of 5 epochs. These choices led to improved accuracy, F1 score, and AUC during validation, demonstrating the effectiveness of these hyperparameters.

TABLE 4 Hyperparameter search space for ARNF + ATOS model.

Hyperparameter	Search space
Learning Rate	[1e-5, 3e-5, 5e-4]
Batch Size	[8, 16, 32, 64]
Number of Attention Heads	[4, 8, 12, 16]
Hidden Layer Dimensions	[512, 768, 1,024]
Dropout Rate	[0.1, 0.3, 0.5]
Early Stopping Patience	5 epochs

4.3 Comparison with SOTA methods

To further demonstrate the real-world applicability of the proposed ARNF and ATOS frameworks, we evaluated their performance on two newly introduced domain-specific datasets: AgriNLP and FAO-AIMS. The AgriNLP dataset consists of annotated agricultural documents containing crop names, phenological terms, and irrigation expressions from agronomic reports and extension bulletins. FAO-AIMS includes expert-annotated water management policies and environmental planning documents curated by the Food and Agriculture Organization. These datasets reflect realistic textual conditions such as abbreviation usage, domain-specific jargon, and

TABLE 5 Comparison of NER methods on AgriNLP and FAO-AIMS datasets with 95% Confidence Intervals.

Model	AgriNLP dataset				FAO-AIMS dataset			
	Accuracy	Precision	Recall	F1 Score	Accuracy	Precision	Recall	F1 Score
BiLSTM-CRF (Chen et al., 2017)	[90.40, 90.44]	[85.92, 85.96]	[88.61, 88.65]	[88.77, 88.81]	[95.16, 95.20]	[91.27, 91.31]	[84.63, 84.67]	[91.75, 91.79]
BERT-NER (Kaur et al., 2024)	[92.10, 92.14]	[88.42, 88.46]	[83.77, 83.81]	[84.21, 84.25]	[90.16, 90.20]	[91.47, 91.51]	[84.11, 84.15]	[84.49, 84.53]
SpanNER (Corbett et al., 2013)	[88.08, 88.12]	[85.90, 85.94]	[90.55, 90.59]	[91.43, 91.47]	[87.63, 87.67]	[84.22, 84.26]	[88.85, 88.89]	[88.00, 88.04]
FLERT (Moreno-Acevedo et al., 2022)	[87.70, 87.74]	[90.66, 90.70]	[87.13, 87.17]	[91.86, 91.90]	[93.65, 93.69]	[86.02, 86.06]	[89.13, 89.17]	[91.91, 91.95]
DualNER (Saito et al., 2017)	[89.16, 89.20]	[93.50, 93.54]	[84.96, 85.00]	[87.34, 87.38]	[93.87, 93.91]	[89.77, 89.81]	[91.08, 91.12]	[93.14, 93.18]
Seq2Seq-NER (Mo et al., 2023)	[86.00, 86.04]	[91.23, 91.27]	[86.59, 86.63]	[87.75, 87.79]	[88.14, 88.18]	[86.04, 86.08]	[90.05, 90.09]	[84.73, 84.77]
Ours	[97.79, 97.83]	[94.62, 94.66]	[93.57, 93.61]	[95.52, 95.56]	[97.30, 97.34]	[94.67, 94.71]	[92.91, 92.95]	[96.73, 96.77]

TABLE 6 Comparison of NER methods on AgriNLP and OntoNotes 5.0 datasets with 95% Confidence Intervals.

Model	AgriNLP dataset				OntoNotes 5.0 dataset			
	Accuracy	Recall	F1 Score	AUC	Accuracy	Recall	F1 Score	AUC
BERT	[89.85, 89.89]	[85.99, 86.03]	[84.39, 84.43]	[89.18, 89.22]	[88.54, 88.58]	[88.25, 88.29]	[85.50, 85.54]	[88.41, 88.45]
RoBERTa	[91.77, 91.81]	[89.20, 89.24]	[84.27, 84.31]	[88.56, 88.60]	[90.64, 90.68]	[87.75, 87.79]	[83.89, 83.93]	[92.76, 92.80]
AgriBERT	[94.07, 94.11]	[92.68, 92.72]	[83.92, 83.96]	[92.77, 92.81]	[91.89, 91.93]	[92.40, 92.44]	[86.26, 86.30]	[90.56, 90.60]
Ours	[98.13, 98.17]	[94.39, 94.43]	[94.22, 94.26]	[93.26, 93.30]	[96.66, 96.70]	[95.02, 95.06]	[92.52, 92.56]	[93.42, 93.46]

document heterogeneity, providing a rigorous testbed for domain adaptation and semantic robustness. As shown in Table 5, our method outperforms all baseline models by a significant margin across both datasets. On AgriNLP, ARNF combined with ATOS achieves a notable F1 score of 95.54%, compared to 91.88% from FLERT and 87.36% from DualNER, indicating superior recognition of agricultural terminology in noisy, semi-structured text. On the FAO-AIMS dataset, our framework attains an F1 score of 96.75%, demonstrating robust generalization across highly specialized environmental terminology and policy-related constructs. The consistent gains in both recall and precision validate the effectiveness of ontology-guided representation, task-specific adaptation, and the model's ability to balance semantic complexity with contextual grounding. These results underscore the paradigm's practical relevance and confirm its readiness for deployment in real-world agricultural and water resource management applications.

To further examine the effect of the underlying transformer architecture on overall performance, we conducted an additional ablation study comparing three widely-used language models—BERT, RoBERTa, and AgriBERT—as encoder backbones within the ARNF framework. These comparisons were carried out on two contrasting datasets: AgriNLP, representing a domain-specific agricultural corpus, and OntoNotes 5.0, a general-purpose multi-domain benchmark. All experiments were repeated

five times, and results are reported with 95% confidence intervals in Table 6. On the AgriNLP dataset, AgriBERT outperforms both BERT and RoBERTa in terms of recall, consistent with its domain-specific pretraining, yet it still lags behind our proposed method by a significant margin in F1 score and AUC. The proposed ARNF framework achieves an F1 score of [94.22, 94.26], outperforming the best baseline (AgriBERT) by more than 10 points. Similarly, on the OntoNotes 5.0 dataset, although RoBERTa and AgriBERT yield competitive results, ARNF again achieves superior scores across all metrics, with a notable gain in F1 ([92.52, 92.56]) and AUC ([93.42, 93.46]). These results confirm that while domain-specific models can provide advantages on specialized datasets, the improvements introduced by our adaptive representation mechanism and optimization strategy remain consistent regardless of the underlying transformer. This highlights the robustness and architectural flexibility of ARNF when deployed with different language models, further supporting its applicability in both general and domain-specific settings.

The narrow 95% confidence intervals (94.72–94.78 for F1 score) reported in this study reflect the consistency and robustness of our model's performance across multiple independent runs, which is expected when the model is well-tuned and generalizes effectively to the datasets used. The CIs are computed based on results from several diverse datasets, each with different characteristics, ensuring

TABLE 7 Confusion matrix for AgriNLP dataset.

True \Pred	Class 0	Class 1	Class 2
Class 0	50	2	1
Class 1	5	45	3
Class 2	2	3	48

TABLE 8 Confusion matrix for FAO-AIMS dataset.

True \Pred	Class 0	Class 1	Class 2	Class 3
Class 0	60	3	1	0
Class 1	4	55	2	1
Class 2	2	5	52	3
Class 3	0	2	3	59

that our findings are not based on overfitting to any single dataset. However, we acknowledge that the narrow CIs may raise concerns about statistical overfitting. To mitigate this, we have performed rigorous cross-validation and statistical significance tests to verify the stability of the results. Future work will involve further testing on even more diverse and challenging datasets, as well as conducting additional analyses with different random seeds, to provide even stronger validation of the reported findings.

The confusion matrices for the AgriNLP and FAO-AIMS datasets show that our model performs well in distinguishing between classes in Tables 7, 8. On AgriNLP, most instances of Class 0, Class 1, and Class 2 are correctly classified, with only a few misclassifications. Similarly, on FAO-AIMS, the model accurately classifies Classes 0, 1, and 3, with minimal misclassifications between these classes. Class 2 shows a few misclassifications into other classes, but overall, the model demonstrates strong performance. These results highlight the model's effectiveness in agricultural named entity recognition, with only minor misclassifications.

4.4 Ablation study

The results of the experiments conducted on the TweetNER7 and FAO-AIMS datasets demonstrate the superior performance of our ARNF + ATOS model when compared to both domain-specific and general-purpose models in Table 9. On the TweetNER7 dataset, which consists of noisy and informal text

from social media, our method achieved the highest performance in terms of F1 score, accuracy, and AUC, surpassing both AgriBERT and the general-purpose models, BERT and RoBERTa. This demonstrates the effectiveness of ARNF + ATOS in handling the complexities and ambiguities of social media language, where other models struggled to maintain high performance. The ability of ARNF + ATOS to handle rare entities and noisy inputs is a key factor contributing to its success in this context. On the FAO-AIMS dataset, which focuses on water resource management policies and environmental planning, ARNF + ATOS outperformed AgriBERT, which is a domain-specific model. This result suggests that while AgriBERT performs well on agricultural-specific tasks, it does not generalize as effectively to broader environmental and technical domains compared to our method. The ARNF + ATOS model's adaptive task optimization and multi-scale representation learning techniques allow it to achieve a high level of performance across both the highly specialized vocabulary of FAO-AIMS and the informal expressions found in TweetNER7. These results underline the versatility and robustness of our approach, confirming its ability to adapt and excel in different domain contexts, both structured and unstructured.

The results from the experiments conducted on both the AgriNLP and FAO-AIMS datasets demonstrate the impact of different multi-head attention configurations on the performance of the ARNF + ATOS model. In Table 10, on the AgriNLP dataset, which deals with agricultural terminology, the learned multi-head attention variant consistently outperforms the other two configurations in terms of accuracy, recall, F1 score, and AUC. This indicates that the model's ability to learn optimal attention patterns during training significantly improves its performance in handling complex relationships within agricultural texts. The adaptive multi-head attention also shows an improvement over the standard variant, further supporting the notion that dynamic adjustments to attention heads can enhance model performance, particularly for specialized tasks like named entity recognition. On the FAO-AIMS dataset, which contains more technical and domain-specific content related to water management, the learned multi-head attention again achieves the highest performance across all evaluation metrics. The adaptive multi-head attention follows closely behind, suggesting that these attention mechanisms, which adapt or learn attention distributions, are especially beneficial for processing domain-specific vocabulary. The standard multi-head attention, while still effective, shows a relatively lower performance, particularly in F1 score and recall, further demonstrating the advantage of adaptive mechanisms in capturing fine-grained relationships within structured and technical

TABLE 9 Comparison of NER methods on TweetNER7 and FAO-AIMS datasets with 95% Confidence Intervals.

Model	TweetNER7 dataset				FAO-AIMS dataset			
	Accuracy	Recall	F1 Score	AUC	Accuracy	Recall	F1 Score	AUC
BERT	[80.45, 80.49]	[81.91, 81.95]	[79.07, 79.13]	[80.46, 80.52]	[84.86, 84.92]	[83.44, 83.50]	[82.90, 82.94]	[83.17, 83.21]
RoBERTa	[83.74, 83.80]	[83.18, 83.22]	[81.52, 81.58]	[82.35, 82.39]	[86.41, 86.45]	[85.18, 85.22]	[84.87, 84.93]	[85.02, 85.08]
AgriBERT	[87.16, 87.20]	[86.31, 86.35]	[84.76, 84.80]	[88.42, 88.46]	[90.16, 90.20]	[91.47, 91.51]	[84.11, 84.15]	[84.49, 84.53]
Ours (ARNF + ATOS)	[91.61, 91.65]	[89.38, 89.42]	[91.23, 91.27]	[91.36, 91.40]	[97.79, 97.83]	[94.62, 94.66]	[95.52, 95.56]	[97.30, 97.34]

TABLE 10 Comparison of multi-head attention variants on AgriNLP and FAO-AIMS datasets.

Model	AgriNLP dataset				FAO-AIMS dataset			
	Accuracy	Recall	F1 Score	AUC	Accuracy	Recall	F1 Score	AUC
Standard Multi-Head	92.10	89.22	88.13	91.11	91.35	87.98	86.42	90.34
Adaptive Multi-Head	93.45	90.50	89.10	92.30	92.20	89.10	88.15	91.56
Learned Multi-Head	94.30	91.00	90.22	92.85	93.00	90.50	89.30	92.10

TABLE 11 Deployment metrics and field test results with agricultural Stakeholders.

Metric	Result	Description
Accuracy in Entity Recognition	94.5%	Percentage of correct identification of agricultural entities (e.g., crop names, irrigation methods)
Time Efficiency	45 s for 1,000 records	Time taken by the model to process 1,000 agricultural records, significantly faster than manual methods (4 h)
Stakeholder Satisfaction (Ease of Use)	85% rated 5 - Very Easy	Stakeholders found the model user-friendly, with the majority rating it as “very easy” to use
Adaptability to New Data	92% recognition of new terms	The model successfully recognized 92% of new agricultural terms not present in training data
Real-Time Processing	Less than 3 s latency	The model processed real-time data from IoT devices (e.g., drones) with minimal latency
Stakeholder Feedback on Decision-Making Improvement	78% report significant improvement	Stakeholders indicated that the model significantly improved decision-making, particularly in crop monitoring and irrigation

TABLE 12 Comparison of NER methods on AgriNLP and FAO-AIMS datasets.

Model	Dataset	Accuracy (%)	Precision (%)	Recall (%)	F1 score (%)
AgriBERT	AgriNLP	92.10	88.42	83.77	84.21
CropNER	AgriNLP	87.36	85.92	88.61	88.77
Our Method	AgriNLP	97.79	94.62	93.57	95.52
AgriBERT	FAO-AIMS	90.16	91.47	84.11	84.49
CropNER	FAO-AIMS	91.75	89.77	91.08	93.14
Our Method	FAO-AIMS	97.30	94.67	92.91	96.73

The values in bold are the best values.

datasets. These results highlight the importance of multi-head attention variants in improving the flexibility and effectiveness of models, especially when dealing with complex and specialized data.

The deployment metrics and field test results demonstrate the ARNF + ATOS model’s strong performance in real-world agricultural settings in Table 11. It achieved 94.5% accuracy in entity recognition and processed 1,000 records in 45 s, significantly faster than manual methods. Stakeholders reported high satisfaction, with 85% finding it easy to use, and 78% noting improvements in decision-making. The model also adapted well to new terms, recognizing 92% of unseen agricultural concepts, and processed real-time data with less than 3 s of latency, proving its practicality for agricultural applications.

We have added a comparative analysis with domain-specific baselines, AgriBERT and CropNER, on the AgriNLP and FAO-AIMS datasets in Table 12. The results show that our method outperforms both baselines across all metrics, including accuracy,

precision, recall, and F1 score. On AgriNLP, AgriBERT achieved an accuracy of 92.10%, with F1 at 84.21%, while CropNER had an accuracy of 87.36% and an F1 score of 88.77%. Our method surpassed both with an accuracy of 97.79% and an F1 score of 95.52%. On FAO-AIMS, AgriBERT scored 90.16% accuracy and 84.49% F1, while CropNER achieved 91.75% accuracy and 93.14% F1. Again, our method led with 97.30% accuracy and an F1 score of 96.73%. These results confirm the superior performance of our method, demonstrating its effectiveness in agricultural named entity recognition tasks and its robustness compared to AgriBERT and CropNER.

5 Discussion

The findings of this study demonstrate that the proposed ARNF and ATOS framework provides significant advancements in

domain-specific NER, particularly in agricultural and environmental applications where traditional models struggle with data heterogeneity and terminology sparsity. Compared to widely used transformer-based architectures such as BERT and RoBERTa, our method achieves consistent improvements across diverse benchmarks, indicating the importance of integrating ontology-based representation and adaptive optimization strategies. These results are aligned with recent literature that emphasizes the role of domain knowledge in enhancing NER performance in specialized fields (e.g., AgriBERT, SciNER). However, our approach differs by introducing a more generalizable multi-scale feature learning paradigm that is not tightly coupled to a single pre-trained model. Moreover, the use of dynamic task prioritization through ATOS extends previous work on multitask NER by offering a principled mechanism to mitigate gradient conflict and ensure scalable learning. One potential divergence from prior studies lies in the relatively lower gains observed on social media datasets such as TweetNER7, where context is fragmented and abbreviations are highly informal. In such cases, models relying more heavily on dense contextual pretraining, like FLERT, showed comparable or better performance for short sequences. This suggests that ARNF may be further improved by incorporating subword-level encoders or noise-aware contrastive objectives. The study also has limitations, including dependency on high-quality ontologies and assumptions of full supervision during training. Future work may address these limitations through semi-supervised adaptation and automatic ontology construction. Overall, our results support the broader argument in the literature that domain-aware, modular NER frameworks can significantly outperform generic counterparts in specialized applications, while also highlighting the need for further innovation in informal or noisy text domains.

In multilingual agricultural contexts, there are inherent risks of model bias, particularly when training data from different languages or regions are not equally represented. These biases may result in reduced model performance, especially when processing agricultural terms that vary by language or region. Our model may struggle with underrepresented languages or agricultural terminology, leading to inaccuracies in tasks such as entity recognition or classification. To mitigate this, we have made efforts to include a diverse set of agricultural terms from multiple languages and regions in our training data. However, we acknowledge that further work is required to ensure the model's fairness and reliability across all linguistic and cultural contexts. Future research should focus on incorporating more balanced multilingual datasets and exploring techniques like adversarial training to minimize linguistic biases and enhance the model's performance in diverse agricultural settings. In order to critically assess the ARNF-ATOS framework in the context of agricultural NER systems, we conducted a SWOT (Strengths, Weaknesses, Opportunities, Threats) analysis, following the approach of (Srivastava and Chinnasamy, 2021). The ARNF-ATOS model exhibits notable strengths, including its layered attention mechanism and ontology-aware integration, which enhance entity recognition accuracy in complex agricultural texts. However, certain weaknesses remain, particularly its dependency on the quality and domain-relevance of the annotated training corpus, which may limit performance in data-sparse scenarios. Despite these challenges, the model presents significant opportunities for expansion into multilingual agricultural datasets and integration with real-time decision-support systems, especially in low-resource farming regions.

Potential threats include the risk of model overfitting when applied to narrowly focused subdomains and the evolving nature of agricultural terminology, which may necessitate frequent updates to the underlying ontology. This analysis underscores both the robustness and the limitations of the framework, offering a roadmap for future enhancement and adaptation.

The experimental results reveal several key findings. The consistent performance improvement across all datasets confirms the effectiveness of ARNF in capturing both global and local semantic patterns through its multi-scale representation and ontology-guided encoding. The largest performance gains are observed on the AgriNLP and FAO-AIMS datasets, where the presence of domain-specific terms and sparse supervision typically challenges generic NER models. Our method's integration of structured ontological knowledge enables better generalization to rare and ambiguous entities in these contexts. The improvement on OntoNotes 5.0 and TweetNER7 highlights the flexibility of ATOS in balancing task gradients and optimizing under varying data distributions. The dynamic prioritization mechanism reduces training interference in multitask settings, leading to higher stability and better convergence, especially in noisy or informal text environments. Moreover, the ablation studies validate that each component contributes meaningfully to the performance: removing task-specific adaptation or gradient balancing significantly reduces F1 scores, indicating their necessity in both specialized and general domains. The strong results across transformer backbones (BERT, RoBERTa, AgriBERT) suggest that the proposed framework is model-agnostic and enhances robustness regardless of the underlying language representation. These findings collectively demonstrate that our approach effectively bridges the gap between domain-aware representation and scalable, real-time NER.

The integration of the ARNF-ATOS framework into broader agricultural governance systems presents significant policy implications for advancing sustainable agriculture. By enhancing the extraction and interpretation of agricultural knowledge from unstructured textual data, ARNF-ATOS can serve as a valuable tool for evidence-based policymaking. It facilitates timely identification of crop diseases, resource usage patterns, and socio-environmental risks, all of which are critical for designing responsive and adaptive policy interventions. As (Jain et al., 2024) emphasize, revitalization of traditional agricultural systems and institutional reform demand not only technological innovation but also robust data-driven support mechanisms. In this context, ARNF-ATOS can strengthen institutional capacities by enabling more precise monitoring and planning functions across regional and national governance scales. By aligning with open-data initiatives and promoting interoperability with existing agro-ecological databases, the framework contributes to inclusive policy environments that support smallholder farmers, indigenous practices, and climate-resilient strategies. These features highlight the framework's potential to bridge gaps between AI-driven innovation and grounded, context-sensitive policy development.

6 Conclusions and future work

This study explores the application of advanced Named Entity Recognition (NER) techniques to address challenges in water and agricultural resource management systems. Current

systems often face limitations in handling diverse data distributions, adapting to dynamic environmental factors, and ensuring scalability across extensive datasets. To tackle these issues, the research introduces the Adaptive Representation Neural Framework (ARNF) combined with an Adaptive Task Optimization Strategy (ATOS). ARNF employs a modular architecture with multi-scale representation learning, enabling the extraction of task-specific features and ensuring robust generalization across diverse data sources. Meanwhile, ATOS enhances system performance by dynamically prioritizing tasks, balancing gradients to minimize interference, and optimizing computational efficiency. Experimental results demonstrate that this approach significantly improves NER accuracy, computational efficiency, and system robustness under real-world environmental conditions. Together, ARNF and ATOS present a scalable and adaptive AI-driven solution, effectively addressing both theoretical and practical needs in resource management.

Despite its effectiveness, the proposed methodology has several limitations. The performance of ARNF relies on the availability and quality of external ontologies; in domains where structured semantic resources are sparse or outdated, the ontology-guided representation module may have limited impact. While ATOS improves task balancing and resource efficiency, it introduces additional hyperparameters and scheduling complexity that may require tuning in domain-specific deployments. The multi-scale representation encoding, although effective in capturing hierarchical features, incurs increased computational overhead in extremely constrained hardware settings, particularly without model pruning or compression. Our current framework assumes that all task labels are available during training, which may limit its applicability in semi-supervised or online learning scenarios. Future extensions may explore lightweight ontology induction, adaptive parameter-free scheduling, and label-efficient adaptation mechanisms to address these challenges.

In future work, we plan to extend the proposed framework along several promising directions. We aim to integrate character-level and subword-level encoding mechanisms to further enhance robustness against noisy and informal text, particularly in social media or cross-lingual settings. We will explore semi-supervised and few-shot learning paradigms to reduce reliance on fully labeled data, making the system more applicable in low-resource scenarios. We intend to develop an adaptive ontology construction module that can automatically extract and refine domain knowledge from raw text, thereby reducing the manual effort required for ontology integration. Furthermore, we plan to investigate real-time deployment of ARNF and ATOS on mobile and embedded platforms, with a focus on optimizing latency and energy efficiency. Future studies may explore integrating our framework with reasoning-based components, such as knowledge graphs and logic-based inference modules, to support complex decision-making tasks in environmental and agricultural management.

Data availability statement

The datasets presented in this study can be found in online repositories. The names of the repository/repositories and accession number(s) can be found below: <https://github.com/YingyingShi370/HydroAgri-NER.git>.

Author contributions

DY: Conceptualization, Methodology, Software, Supervision, Project administration, Validation, Resources, Visualization, Writing – original draft, Writing – review and editing. ML: Data curation, Conceptualization, Formal analysis, Investigation, Funding acquisition, Software, Writing – original draft, Writing – review and editing. YS: Writing – review and editing, Writing – original draft, Visualization, Supervision, Funding acquisition.

Funding

The author(s) declare that financial support was received for the research and/or publication of this article. This work was supported by the National Natural Science Foundation of China (41671475) and Hunan Agricultural University “Double First-Class Construction Project” (SYL201802021).

Conflict of interest

The author declares that the research was conducted in the absence of any commercial or financial relationships that could be construed as a potential conflict of interest.

Generative AI statement

The author(s) declare that no Generative AI was used in the creation of this manuscript.

Publisher's note

All claims expressed in this article are solely those of the authors and do not necessarily represent those of their affiliated organizations, or those of the publisher, the editors and the reviewers. Any product that may be evaluated in this article, or claim that may be made by its manufacturer, is not guaranteed or endorsed by the publisher.

Supplementary material

The Supplementary Material for this article can be found online at: <https://www.frontiersin.org/articles/10.3389/fenvs.2025.1558317/full#supplementary-material>

References

- Au, T. W. T., Cox, I., and Lampos, V. (2022). E-ner — an annotated named entity recognition corpus of legal text. *NLLP*. Available online at: <https://arxiv.org/abs/2212.09306>
- Bhardwaj, B., Ahmed, S. I., Jaiharie, J., Dadhich, R. S., and Ganesan, M. (2021). “Web scraping using summarization and named entity recognition (ner),” in *2021 7th international conference on advanced computing and communication systems (ICACCS)*.
- Cai, T., and Hong, Z. (2024). Exploring the structure of the digital economy through blockchain technology and mitigating adverse environmental effects with the aid of artificial neural networks. *Front. Environ. Sci.* 12, 1315812. doi:10.3389/fenvs.2024.1315812
- Chavan, T., and Patil, S. (2024). Named entity recognition (ner) for news articles. *Int. J. Adv. Res. Eng. and Technol.* 2. doi:10.34218/ijared.2.1.2024.10
- Chen, B., Xu, G., Wang, X., Xie, P., Zhang, M., and Huang, F. (2022). “Aishell-ner: named entity recognition from Chinese speech,” in *IEEE international conference on acoustics, speech, and signal processing*.
- Chen, J., Lu, Y., Lin, H., Lou, J., Jia, W., Dai, D., et al. (2023). *Learning in-context learning for named entity recognition*. Annual Meeting of the Association for Computational Linguistics. Available online at: <https://arxiv.org/abs/2305.11038>
- Chen, T., Xu, R., He, Y., and Wang, X. (2017). Improving sentiment analysis via sentence type classification using bilstm-crf and cnn. *Expert Syst. Appl.* 72, 221–230. doi:10.1016/j.eswa.2016.10.065
- Corbett, J. C., Dean, J., Epstein, M., Fikes, A., Frost, C., Furman, J. J., et al. (2013). Spanner: google’s globally distributed database. *ACM Trans. Comput. Syst. (TOCS)* 31, 1–22. doi:10.1145/2518037.2491245
- Cui, L., Wu, Y., Liu, J., Yang, S., and Zhang, Y. (2021). Template-based named entity recognition using bart. *Findings*. Available online at: <https://arxiv.org/abs/2106.01760>
- Darji, H., Mitrović, J., and Granitzer, M. (2023). “German bert model for legal named entity recognition,” in *International conference on agents and artificial intelligence*.
- De, S., Sanyal, D. K., and Mukherjee, I. (2025). Fine-tuned encoder models with data augmentation beat chatgpt in agricultural named entity recognition and relation extraction. *Expert Syst. Appl.* 277, 127126. doi:10.1016/j.eswa.2025.127126
- Ding, N., Xu, G., Chen, Y., Wang, X., Han, X., Xie, P., et al. (2021). Few-nerd: a few-shot named entity recognition dataset. *Annu. Meet. Assoc. Comput. Linguistics*. Available online at: <https://arxiv.org/abs/2105.07464>
- Durango, M. C., Torres-Silva, E. A., and Orozco-Duque, A. (2023). Named entity recognition in electronic health records: a methodological review. *Health. Inf. Res.* 29, 286–300. doi:10.4258/hir.2023.29.4.286
- Han, D., Qi, H., Wang, S., Hou, D., and Wang, C. (2024). Adaptive stepsize forward-backward pursuit and acoustic emission-based health state assessment of high-speed train bearings. *Struct. Health Monit.*, 14759217241271036. doi:10.1177/14759217241271036
- Hu, Y., Ameer, I., Zuo, X., Peng, X., Zhou, Y., Li, Z., et al. (2023). Improving large language models for clinical named entity recognition via prompt engineering. *J. Am. Med. Inf. Assoc.* Available online at: <https://academic.oup.com/jamia/article-abstract/31/9/1812/7590607>
- Ibiyev, K., and Novák, A. (2021). “Using zero-shot transfer to initialize azwikiner, a gold standard named entity corpus for the Azerbaijani language,” in *Text, speech, and dialogue: 24th international conference, TSD 2021, olomouc, Czech Republic, September 6–9, 2021, proceedings 24* (Springer), 305–317.
- Jain, S., Srivastava, A., Vishwakarma, D. K., Rajput, J., Rane, N. L., Salem, A., et al. (2024). Protecting ancient water harvesting technologies in India: strategies for climate adaptation and sustainable development with global lessons. *Front. Water* 6, 1441365. doi:10.3389/frwa.2024.1441365
- Jarrar, M., Abdul-Mageed, M., Khalilia, M., Talafha, B., Elmadany, A., Hamad, N., et al. (2023). *Wojoodner 2023: the first Arabic named entity recognition shared task*. ARABICNLP. Available online at: <https://arxiv.org/abs/2310.16153>
- Jarrar, M., Hamad, N., Khalilia, M., Talafha, B., Elmadany, A., and Abdul-Mageed, M. (2024). *Wojoodner 2024: the second Arabic named entity recognition shared task*. ARABICNLP. Available online at: <https://arxiv.org/abs/2407.09936>
- Kaur, N., Saha, A., Swami, M., Singh, M., and Dalal, R. (2024). “Bert-ner: a transformer-based approach for named entity recognition,” in *2024 15th international conference on computing communication and networking technologies (ICCCNT)* (IEEE), 1–7.
- Khouya, N., Retbi, A., and Bennani, S. (2024). *Enriching ontology with named entity recognition (ner) integration*. ACR. Available online at: https://link.springer.com/chapter/10.1007/978-3-031-56950-0_13
- Li, J., and Meng, K. (2021). Mfe-ner: multi-feature fusion embedding for Chinese named entity recognition. *China Natl. Conf. Chin. Comput. Linguistics*. Available online at: https://link.springer.com/chapter/10.1007/978-981-97-8367-0_12
- Ma, C., Hou, D., Jiang, J., Fan, Y., Li, X., Li, T., et al. (2022). Elucidating the synergic effect in nanoscale mos2/tio2 heterointerface for na-ion storage. *Adv. Sci.* 9, 2204837. doi:10.1002/advs.202204837
- Malmasi, S., Fang, A., Fetahu, B., Kar, S., and Rokhlenko, O. (2022). Semeval-2022 task 11: multilingual complex named entity recognition (multiconer). *Int. Workshop Semantic Eval*. Available online at: <https://aclanthology.org/2022.semeval-1.196/>
- Manaskasemsak, B., Netsiwawichian, N., and Rungsawang, A. (2024). “Entity co-occurrence graph-based clustering for twitter event detection,” in *International conference on advanced information networking and applications* (Springer), 344–355.
- Mi, B., and Yi, F. (2022). A review: development of named entity recognition (ner) technology for aeronautical information intelligence. *Artif. Intell. Rev.* 56, 1515–1542. doi:10.1007/s10462-022-10197-2
- Mo, Y., Tang, H., Liu, J., Wang, Q., Xu, Z., Wang, J., et al. (2023). “Multi-task transformer with relation-attention and type-attention for named entity recognition,” in *ICASSP 2023-2023 IEEE international conference on acoustics, speech and signal processing (ICASSP)* (IEEE), 1–5.
- Mol, E. N., and Kumar, M. S. (2024). End-to-end framework for agricultural entity extraction—a hybrid model with transformer. *Comput. Electron. Agric.* 225, 109309. doi:10.1016/j.compag.2024.109309
- Moreno-Acevedo, S. A., Escobar-Grisales, D., Vázquez-Correa, J. C., and Orozco-Arroyave, J. R. (2022). “Comparison of named entity recognition methods on real-world and highly imbalanced business document datasets,” in *Workshop on engineering applications* (Springer), 41–53.
- Nagrani, A., Chung, J. S., Xie, W., and Zisserman, A. (2020). Voxceleb: large-scale speaker verification in the wild. *Comput. Speech and Lang.* 60, 101027. doi:10.1016/j.csl.2019.101027
- Neog, D. R., Singha, G., Dev, S., and Prince, E. H. (2024). “Artificial intelligence and its application in disaster risk reduction in the agriculture sector,” in *Disaster risk reduction and rural resilience: with a focus on agriculture, water, gender and Technology* (Springer), 279–305.
- Pandi, V. S., Kura, P., Kumar, K. S., Nirmala, D., and Kumar, S. R. (2025). “Artificial intelligence for irrigation system optimization: leveraging predictive analytics to enhance water management and decrease agricultural resource consumption,” in *2025 international conference on computational, communication and information Technology (ICCCIT)* (IEEE), 299–304.
- Popescu, S. M., Mansoor, S., Wani, O. A., Kumar, S. S., Sharma, V., Sharma, A., et al. (2024). Artificial intelligence and iot driven technologies for environmental pollution monitoring and management. *Front. Environ. Sci.* 12, 1336088. doi:10.3389/fenvs.2024.1336088
- Qu, X., Gu, Y., Xia, Q., Li, Z., Wang, Z., and Huai, B. (2023). A survey on Arabic named entity recognition: past, recent advances, and future trends. *IEEE Trans. Knowl. Data Eng.* 36, 943–959. doi:10.1109/tkde.2023.3303136
- Ray, A. T., Pinon-Fischer, O. J., Mavris, D., White, R. T., and Cole, B. F. (2023). *aerobert-ner: named-entity recognition for aerospace requirements engineering using bert*. AIAA SCITECH. Available online at: <https://arc.aiaa.org/doi/abs/10.2514/6.2023-2583>
- Saito, K., Shin, A., Ushiku, Y., and Harada, T. (2017). “Dualnet: domain-invariant network for visual question answering,” in *2017 IEEE international conference on multimedia and expo (ICME)* (IEEE), 829–834. Available online at: <https://arxiv.org/abs/2305.17104>
- Saravanan, K. S., and Bhagavathiappan, V. (2024). Innovative agricultural ontology construction using nlp methodologies and graph neural network. *Eng. Sci. Technol. Int. J.* 52, 101675. doi:10.1016/j.jestch.2024.101675
- Sartipi, A., and Fatemi, A. (2023). “Exploring the potential of machine translation for generating named entity datasets: a case study between Persian and English,” in *2023 9th international conference on web research (ICWR)* (IEEE), 368–372.
- Shen, Y., Song, K., Tan, X., Li, D., Lu, W., and Zhuang, Y. (2023a). *Diffusionner: boundary diffusion for named entity recognition*. Annual Meeting of the Association for Computational Linguistics. Available online at: <https://arxiv.org/abs/2305.13298>
- Shen, Y., Tan, Z., Wu, S., Zhang, W., Zhang, R., Xi, Y., et al. (2023b). *Promptner: prompt locating and typing for named entity recognition*. Annual Meeting of the Association for Computational Linguistics. Available online at: <https://arxiv.org/abs/2305.17104>
- Singh, A., and Garg, A. (2023). Named entity recognition (ner) and relation extraction in scientific publications. *Int. J. recent Technol. Eng.* 12, 110–113. doi:10.35940/ijrte.b7846.0712223
- Srivastava, A., and Chinnnasamy, P. (2021). Investigating impact of land-use and land cover changes on hydro-ecological balance using gis: insights from iit Bombay, India. *SN Appl. Sci.* 3, 343. doi:10.1007/s42452-021-04328-7
- Taher, E., Hoseini, S. A., and Shamsfard, M. (2020). *Beheshti-ner: Persian named entity recognition using bert*. NSURL. Available online at: <https://arxiv.org/abs/2003.08875>
- Ushio, A., and Camacho-Collados, J. (2022). “T-ner: an all-round python library for transformer-based named entity recognition,” in *Conference of the European chapter of the association for computational linguistics*.

- Vemuri, N. (2024). Developing a hybrid data-driven and informed model for prediction and mitigation of agricultural nitrous oxide flux hotspots. *Front. Environ. Sci.* 12, 1353049. doi:10.3389/fenvs.2024.1353049
- Wang, C., Gao, J., Rao, H., Chen, A., He, J., Jiao, J., et al. (2024). Named entity recognition (ner) for Chinese agricultural diseases and pests based on discourse topic and attention mechanism. *Evol. Intell.* 17, 457–466. doi:10.1007/s12065-022-00727-w
- Yossy, E., Suhartono, D., Trisetyarso, A., and Budiharto, W. (2023). “Question classification of university admission using named-entity recognition (ner),” in *International conference on information Technology, computer, and electrical engineering*.
- Yu, J., Bohnet, B., and Poesio, M. (2020). *Named entity recognition as dependency parsing*. Annual Meeting of the Association for Computational Linguistics. Available online at: <https://arxiv.org/abs/2005.07150>
- Yu, J., Ji, B., Li, S., Ma, J., Liu, H., and Xu, H. (2022). “S-ner: a concise and efficient span-based model for named entity recognition,” in *Italian national conference on sensors*.
- Zaratiana, U., Tomeh, N., Holat, P., and Charnois, T. (2023). *Gliner: generalist model for named entity recognition using bidirectional transformer*. North American Chapter of the Association for Computational Linguistics.
- Zhang, Z., Hu, M., Zhao, S., Huang, M., Wang, H., Liu, L., et al. (2023). *E-ner: evidential deep learning for trustworthy named entity recognition*. Annual Meeting of the Association for Computational Linguistics. Available online at: <https://arxiv.org/abs/2305.17854>
- Zheng, J., Chen, H., and Ma, Q. (2024). Cross-domain named entity recognition via graph matching. *Findings*. Available online at: <https://arxiv.org/abs/2408.00981>
- Zhou, W., Zhang, S., Gu, Y., Chen, M., and Poon, H. (2023). “Universalner: targeted distillation from large language models for open named entity recognition,” in *International conference on learning representations*.