

# Ellipsoid-Based Decision Boundaries for Open Intent Classification

Yuetian Zou<sup>1\*</sup>, Hanlei Zhang<sup>1\*</sup>, Hua Xu<sup>1†</sup>, Songze Li<sup>2</sup>, Long Xiao<sup>2</sup>

<sup>1</sup>Tsinghua University

<sup>2</sup>Hebei University of Science and Technology

zouyt21@mails.tsinghua.edu.cn, zhang-hl20@mails.tsinghua.edu.cn, xuhua@mail.tsinghua.edu.cn

## Abstract

Textual open intent classification is crucial for real-world dialogue systems, enabling robust detection of unknown user intents without prior knowledge and contributing to the robustness of the system. While adaptive decision boundary methods have shown great potential by eliminating manual threshold tuning, existing approaches assume isotropic distributions of known classes, restricting boundaries to balls and overlooking distributional variance along different directions. To address this limitation, we propose EliDecide, a novel method that learns ellipsoid decision boundaries with varying scales along different feature directions. First, we employ supervised contrastive learning to obtain a discriminative feature space for known samples. Second, we apply learnable matrices to parameterize ellipsoids as the boundaries of each known class, offering greater flexibility than spherical boundaries defined solely by centers and radii. Third, we optimize the boundaries via a novel designed dual loss function that balances empirical and open-space risks: expanding boundaries to cover known samples while contracting them against synthesized pseudo-open samples. Our method achieves state-of-the-art performance on multiple text intent benchmarks and further on a question classification dataset. The flexibility of the ellipsoids demonstrates superior open intent detection capability and strong potential for generalization to more text classification tasks in diverse complex open-world scenarios.

**Code** — <https://github.com/thuiar/textoir>

**Extended version** — <https://arxiv.org/abs/2511.16685>

## Introduction

Open world classification is critical for robust real-world systems, as it enables accurate recognition of known classes while simultaneously rejecting anomalies to reduce open-space risks (Yang et al. 2024b). This capability is essential across diverse domains, including autonomous driving (Geiger, Lenz, and Urtasun 2012) and medical image analysis (Zimmerer et al. 2022). In the realm of natural language understanding, dialogue systems specifically require

\*These authors contributed equally.

†This is the corresponding author.

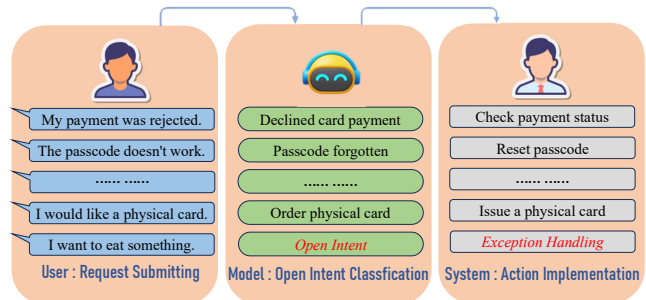


Figure 1: Open Intent Classification in Banking: Model classifies user requests into known intents (subsequently triggering corresponding services), or detects unknown (*open*) intents (handled as exceptions).

the ability to recognize the known user intents while detecting unknown ones, which is a significant capability for system robustness and continuous improvement (Lin and Xu 2019), as shown in Figure 1. Consequently, to develop robust dialogue systems, extensive research (Lin and Xu 2019; Larson et al. 2019; Zhang, Xu, and Lin 2021; Zhang et al. 2023a; Zhou, Liu, and Qiu 2022; Zhou et al. 2023) has focused on open world classification for user intents, known as *Open Intent Classification*. In this work, we formulate open world classification as a  $(K+1)$ -way classification task comprising  $K$  known classes and a single unknown (*open*) class. This formulation aligns with the established paradigm of Open Set Recognition (OSR) (Scheirer et al. 2012).

While the early open world classification methods (Bendale and Boult 2016; Hendrycks and Gimpel 2018; Liang, Li, and Srikant 2020) relied on class probability distributions, recent approaches utilize deep representations from pretrained models, which primarily fall into two categories. Scoring-based approaches (Zhou, Liu, and Qiu 2022; Zhou et al. 2023; Zhang et al. 2023b; Yang et al. 2024a; Gautam et al. 2024) construct discriminative feature spaces and identify unknown samples through statistic-based scoring functions like local density measures, but they require manual threshold tuning, limiting their practical use. In contrast, boundary-based methods (Zhang, Xu, and Lin 2021; Zhang et al. 2023a; Liu et al. 2023; Chen et al. 2024; Li et al. 2025) avoid thresholds by learning adaptive decision boundaries.




Figure 2: An example of a known class’s anisotropic feature distribution. The ball boundary excludes a significant portion of known samples as the cost of avoiding open ones inside, while the ellipsoid includes most of the known samples without many open ones inside due to geometric flexibility.

However, current boundary-based approaches typically assume isotropic feature distributions, ignoring the directional variance inherent in real-world data. As illustrated in Figure 2, ellipsoid boundaries offer a more flexible representation than spherical ones, yielding lower open-space and empirical risks while producing more concise closed regions.

In this paper, we propose EliDecide, an ellipsoid-based decision boundary learning method for open-world classification. First, supervised contrastive learning is applied to a pre-trained model to obtain discriminative representations that serve as the basis for adaptive boundary construction. In the feature representation space, we construct ellipsoid-shaped boundaries for each known class and regard the outside of all ellipsoids as the open space. Each ellipsoid is parameterized by a learnable nonsingular matrix encoding both the directions and the lengths of the ellipsoid’s axes. We optimize the boundaries using a dual loss mechanism. A positive loss expands the boundary when known samples (used as positive samples) fall outside it to encourage it to involve more known samples, while a negative loss contracts the boundary when pseudo-open samples (used as negative examples) lie too close or within it to prevent it from over-expanding and involving unknown samples. This integrated strategy enables EliDecide to adapt decision boundaries to the inherent directional variance in real-world data, representing the distributions of known samples flexibly and precisely and achieving robust open-world classification.

We extensively evaluate EliDecide on benchmark intent datasets and demonstrate its superior performance and robustness compared to state-of-the-art methods. Beyond intent, experiments on a question classification dataset show the generalization potential of our method in more diverse and complex open-world classification tasks. Additional experiments including comparisons among boundaries of different shapes further validate the advantage of the ellipsoid formulation, highlighting its effectiveness and robustness across different open-world scenarios.

We summarize our contributions as follows:

- To overcome the limitation of spherical boundaries in existing methods, we introduce the first ellipsoid boundary method for open world classification. We propose an effective parameterization that represents an ellipsoid via a learnable matrix capturing both axis directions and lengths. Novel and effective dual loss mechanism is de-

signed to optimize the parameter matrices of boundaries.

- We evaluate EliDecide on two benchmark intent datasets as well as a question classification dataset, achieving a new state-of-the-art performance across all tasks. This demonstrates its robust effectiveness and strong generalization capabilities.
- Comprehensive additional experiments and analysis including geometric interpretation, ablation studies and performance comparisons with Large Language Models (LLMs) validate the theoretical foundations and practical superiority of our methodology.

## Related Works

Based on the primary detection criterion, existing methods fall into two categories: probability-based methods and deep representation-based methods.

### Probability-Based Methods

Early-stage methods detect open samples using class probability distributions. For example, OpenMax (Bendale and Boulton 2016) modifies softmax networks by replacing the penultimate layer with activation vectors to obtain the probabilities of  $(K + 1)$ -way classification task. Softmax-MSP (Hendrycks and Gimpel 2018) demonstrates that softmax probabilities can serve as a baseline for distinguishing between known and unknown data, while ODIN (Liang, Li, and Srikant 2020) improves this approach by applying temperature scaling and input perturbations. However, these methods rely solely on output probabilities and fail to leverage the deep semantic features.

### Deep Representation-Based Methods

Recent approaches overcome the limitations of softmax-based classifiers by exploiting deep semantic features extracted by pre-trained models. These methods can be further categorized as *scoring-based methods* and *adaptive boundary-based methods*.

**Scoring-based methods** learn discriminative feature representations to calculate confidence scores and discriminate between known and open samples by comparing the scores against a threshold. For example, DOC (Shu, Xu, and Liu 2017) employs a one-vs-rest sigmoid layer to detect open samples. Lin and Xu (2019) first explores open intent detection using bidirectional LSTMs with local outlier factor (LOF) techniques. KNNCL (Zhou, Liu, and Qiu 2022) uses contrastive learning to improve the discriminability of semantic features. DE (Zhou et al. 2023) addresses deep over-confidence by voting based-early exiting. SCOOS (Gautam et al. 2024) leverages the embeddings of the intent labels as semantic cues. Both KNNCL and DE adopt the density-based scoring functions like LOF for detecting open samples while SCOOS uses statistic-based thresholds. However, manual threshold tuning hinders their practical deployment.

**Adaptive boundary-based methods** dynamically adjust decision boundaries instead of relying on manually tuned thresholds. Zhang, Xu, and Lin (2021) further formulates the open intent classification task and proposes ball-like adaptive decision boundaries to reduce open-space risk. On this

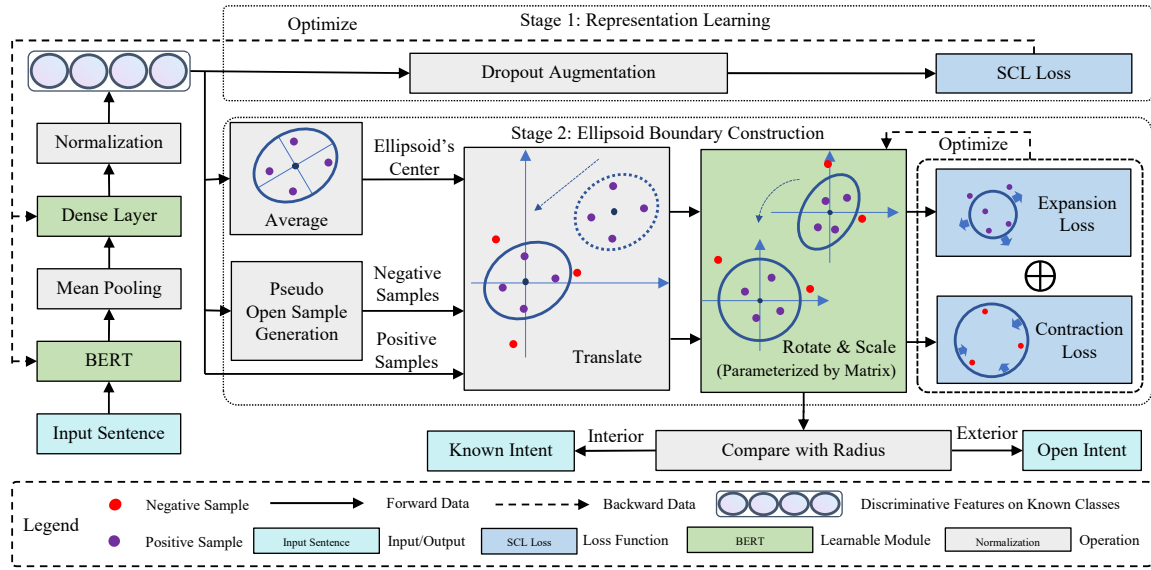


Figure 3: The structure of our method. The first stage is to learn discriminative representations by SCL and then learn ellipsoid-based boundaries with expansion and contraction losses by using both positive and pseudo-negative samples.

basis, to refine the feature representation for the boundaries, DA-ADB (Zhang et al. 2023a) introduces Distance-Aware Representation Learning and TCAB (Chen et al. 2024) applies a triplet-contrastive learning strategy. CLAB (Liu et al. 2023) designs a novel loss function with expanding and shrinking operations to learn more suitable boundaries. MOGB (Li et al. 2025) establishes multiple spherical boundaries for individual classes. Different from the existing methods with spherical boundaries, our proposed method introduces a more flexible and adaptive ellipsoidal boundaries that provide discriminative criteria for this challenging task.

## Proposed Method

Our proposed method aims to build an ellipsoid-shaped boundary to describe the distribution of each known class in the feature space. The complementary space to all the ellipsoids will be regarded as the open space. Our method can be separated into two stages as shown in Figure 3. In the first stage, we apply supervised contrastive learning (Khosla et al. 2020) to learn discriminative representations for boundary learning. Then, we construct ellipsoid-based boundaries and optimize them with novelly designed dual optimization objectives by using both positive and pseudo-negative samples.

### Representation Learning

Following (Zhang, Xu, and Lin 2021; Zhou, Liu, and Qiu 2022; Zhou et al. 2023), we apply pre-trained BERT model (Devlin et al. 2019) as the backbone for feature extraction and use the average representations of the last layer of the encoder to yield the sentence-level representation:

$$x_i = \frac{1}{N+1}([\text{CLS}] + \sum_{j=1}^N T_j), z_i = \frac{Wx_i + b}{\|Wx_i + b\|_2}, \quad (1)$$

where  $N$  is the number of tokens in the sentence,  $[\text{CLS}]$  and  $\{T_j\}_{j=1}^N \in \mathbb{R}^H$  are token-level features,  $H$  is the hidden size. As suggested in (Zhang et al. 2023a), we add a linear layer with the weight matrix  $W \in \mathbb{R}^{H \times n}$  and bias vector  $b \in \mathbb{R}^n$  to enhance the vanilla representation  $x_i$ , where  $n$  is the dimension of the representation. Normalization is applied to ensure compactness of the representation space and harmonize the scales of the distributions among classes, yielding the final representation  $z_i \in \mathbb{R}^n$ .

With the motivation of pushing intra-class compact and inter-class separation, we apply supervised contrastive learning (Khosla et al. 2020) as below:

$$\mathcal{L}_{\text{con}} = - \sum_{i \in I} \frac{1}{|P_i|} \sum_{p \in P_i} \log \frac{\exp(z_i \cdot z_p / \tau)}{\sum_{a \in A_i} \exp(z_i \cdot z_a / \tau)}, \quad (2)$$

where  $I$  are the sample indices in a mini-batch,  $A_i = I \setminus \{i\}$  and  $P_i = \{p \in A(i) | y_p = y_i\}$ . Here we apply *dropout augmentation* (Gao, Yao, and Chen 2021) to obtain the positive augmentation of each sample. By minimizing  $\mathcal{L}_{\text{con}}$ , the representations of samples with the same label are encouraged to converge, whereas those of samples with distinct labels are induced to diverge.

### Ellipsoid Boundary Construction

For each known class  $k$ , we denote the region inside its ellipsoid boundary as  $E_k \subset \mathbb{R}^n$ . To construct this boundary, we first compute the centroid of the ellipsoid and then learn a matrix that encodes its *shape*, including both its size and orientation. Simultaneously, we construct an affine transformation that maps the ellipsoid to a ball, which is useful for determining whether a point lies within the ellipsoid and for calculating the proposed loss function.

**Centroid Computation** For each known class  $k$ , we compute the ellipsoid centroid  $c_k$  as the mean of the representa-

tions of samples belonging to class  $k$ :

$$c_k = \frac{1}{|S_k|} \sum_{z_i \in S_k} z_i, \quad (3)$$

where  $S_k = \{z_i \mid y_i = k\}$  denotes the representation set of class  $k$ .

**Shape Parameterization via Learnable Matrix** In this section we parameterize an ellipsoid  $\widetilde{E}_k$  with the center at the origin, namely a *centralized ellipsoid*, as a pattern of the shape (size and orientation) of  $E_k$ . Then  $E_k$  is defined by a simple translation:

$$E_k = \widetilde{E}_k + c_k := \{\hat{x} \in \mathbb{R}^n \mid \hat{x} = x + c_k, x \in \widetilde{E}_k\}. \quad (4)$$

A centralized ellipsoid is commonly characterized by  $n$  orthogonal unit vectors  $\{d_i\}_{i=1}^n$  specifying its axis directions and corresponding positive axis lengths  $\{a_i\}_{i=1}^n \subset \mathbb{R}^+$ . However, directly optimizing the orthogonal vectors  $\{d_i\}$  using gradient descent is intractable due to the rigid orthogonality constraints. To overcome this, we propose to learn a matrix  $A_k \in \mathbb{R}^{n \times n}$  that implicitly encodes both the axis directions and scales and define the centralized ellipsoid as

$$\widetilde{E}_k = \{x \in \mathbb{R}^n \mid \|A_k x\|_2 \leq \Delta_k\}, \quad (5)$$

where  $\Delta_k = \frac{1}{|S_k|} \sum_{z_i \in S_k} \|z_i - c_k\|_2$  is the average distance between the centroid  $c_k$  and samples with label  $k$  to enhance the numerical stability among known classes with different scales. This formulation eliminates the need for explicit orthogonality constraints while ensuring that  $\widetilde{E}_k$  remains an ellipsoid.

To demonstrate that  $\widetilde{E}_k$  defined by Eq.5 is indeed an ellipsoid, we start by expanding the quadratic form induced by  $A_k$  to reveal the structure of the ellipsoid, which then leads naturally to the construction of a symmetric matrix and its spectral interpretation.

**Quadratic Form Expansion:** Let  $A_k = [a_{m,i}] \in \mathbb{R}^{n \times n}$  be the learnable matrix and  $x = (x_1, \dots, x_n)^\top \in \mathbb{R}^n$ . Then the squared norm  $\|A_k x\|_2^2$  can be expanded as:

$$\begin{aligned} \|A_k x\|_2^2 &= \sum_{i=1}^n \left( \sum_{k=1}^n a_{k,i} \right)^2 x_i^2 \\ &+ 2 \sum_{1 \leq i < j \leq n} \left( \sum_{k=1}^n a_{k,i} a_{k,j} \right) x_i x_j, \end{aligned} \quad (6)$$

which reveals both the quadratic and cross terms.

**Symmetric Matrix Construction:** We define a symmetric matrix  $\widehat{A}_k = [\widehat{a}_{i,j}]$  by

$$\begin{aligned} \widehat{a}_{i,i} &= \sum_{m=1}^n a_{m,i}^2 \quad (\text{diagonal terms}), \\ \widehat{a}_{i,j} &= \widehat{a}_{j,i} = \sum_{m=1}^n a_{m,i} a_{m,j} \quad (i \neq j), \end{aligned} \quad (7)$$

which exactly captures the coefficients appearing in (6). Therefore, we can rewrite

$$\|A_k x\|_2^2 = x^\top \widehat{A}_k x. \quad (8)$$

**Spectral Interpretation:** The centralized ellipsoid can be equivalently expressed as

$$\widetilde{E}_k = \{x \in \mathbb{R}^n \mid x^\top \widehat{A}_k x \leq \Delta_k^2\}. \quad (9)$$

Since  $\widehat{A}_k$  is symmetric positive-definite, its spectral decomposition is given by  $\widehat{A}_k = \sum_{i=1}^n \lambda_i d_i d_i^\top$ , where  $\{d_i\}_{i=1}^n$  are orthonormal eigenvectors and  $\lambda_i > 0$  are the corresponding eigenvalues. This implies that the ellipsoid's axes are aligned with the eigenvectors  $d_1, \dots, d_n$  and have lengths  $a_i = \frac{\Delta_k}{\sqrt{\lambda_i}}$ .

It has been proved that  $\widetilde{E}_k$  is a centralized ellipsoid whose axes are  $\{d_i\}_{i=1}^n$  and their lengths are  $\{a_i\}_{i=1}^n$  correspondingly, and thereby  $E_k$  is also an ellipsoid parameterized as:

$$E_k = \{x \in \mathbb{R}^n \mid \|A_k(x - c_k)\|_2 \leq \Delta_k\} \quad (10)$$

by (4) and (5).

Finally, we define the affine transformation

$$\varphi_k(z) = A_k(z - c_k) \quad (11)$$

that map  $E_k$  to a centralized ball. In particular,

$$z \in E_k \iff \|\varphi_k(z)\|_2 \leq \Delta_k. \quad (12)$$

This equivalence provides a convenient parametrization for points within the ellipsoid.

**Boundary Optimization Objectives** To establish discriminative class boundaries, we design dual-directional boundary adjustment mechanisms through complementary loss functions. For known samples, boundaries are optimized to expand and encompass their distribution, while synthesized pseudo-negative samples induce boundary contraction in under-explored regions.

**Expansion Loss:** Given a labeled sample  $z$  from class  $k$ , we define the expansion loss as:

$$\mathcal{L}_{p,k}(z) := \max(\|\varphi_k(z)\|_2 - \Delta_k, 0) \quad (13)$$

This loss activates when  $z$  resides outside the current class boundary ( $\|\varphi_k(z)\|_2 > \Delta_k$ ), creating gradient signals that *expand* the boundary radially until it encloses the sample. Continuous application of this mechanism enables dynamic boundary adaptation to the inherent data distribution.

**Contraction Loss:** To mitigate over-generalization into regions without known samples, we generate pseudo-open samples as negative samples to encourage the boundaries to contract in those areas. The pseudo-open samples are synthesized by mixing  $P$  known samples with distinct labels (Zhang et al. 2024). Specifically, a pseudo-open sample is computed as

$$z' = \sum_{i=1}^P \lambda_i z_i, \quad \text{with } \lambda \sim \text{Dir}(\alpha), \quad (14)$$

where  $P$  and  $\alpha$  are hyperparameters, and  $\{z_i\}_{i=1}^P$  are known samples with  $P$  different labels. The coefficients  $\{\lambda_i\}$  are drawn from a Dirichlet distribution (ensuring  $\sum_{i=1}^P \lambda_i = 1$  and  $\lambda_i \in [0, 1]$ ). Choosing  $\alpha < 1$  typically forces one coefficient to dominate, so that each pseudo-open sample tends to lie close to one known class while still exhibiting slight influences from the others.

The motivation behind the contraction loss is to push the boundary inward when a negative sample lies too close to or inside the boundary. For a negative sample  $z'$  with respect to class  $k$ , the negative loss is defined as:

$$\mathcal{L}_{n,k}(z') = \begin{cases} (\Delta_k - r_k(z')) + \beta, & r_k(z') < \Delta_k \\ \beta e^{\Delta_k - r_k(z')}, & r_k(z') \geq \Delta_k \end{cases} \quad (15)$$

where  $r_k(z') = \|\varphi_k(z')\|_2$ ,  $\beta$  is a hyperparameter that controls the penalty strength. Specifically, when  $\|\varphi_k(z')\|_2 < \Delta_k$ , the loss increases linearly as the negative sample penetrates deeper into the boundary, thus exerting a direct contraction force. When  $\|\varphi_k(z')\|_2 \geq \Delta_k$ , the exponential term applies a softer penalty, ensuring a gradual enforcement of the boundary’s inward contraction. The parameter  $\beta$  effectively balances these two regimes, with larger values of  $\beta$  inducing a stronger contraction effect on the boundary.

The final optimization objective  $\mathcal{L}_{\text{tot}}$  aggregates per-sample losses across all classes:

$$\mathcal{L}_{\text{tot}} = \sum_{k=1}^K \left[ \sum_{z \in \mathcal{B} \cap S_k} \mathcal{L}_{p,k}(z) + \sum_{z' \in \mathcal{B}'} \mathcal{L}_{n,k}(z') \right] \quad (16)$$

where  $\mathcal{B}$  denotes a training batch and  $\mathcal{B}'$  is a corresponding set of negative samples of the same size, generated by  $\mathcal{B}$ . This formulation jointly coordinates boundary expansion and contraction through gradient signals from both positive and negative samples.

## Inference

For a given sample  $z$ , we first assign it to the nearest class based on the Euclidean distance to class centroids:

$$j = \arg \min_{k \in \{1, \dots, K\}} \|z - c_k\|_2. \quad (17)$$

Next, we check whether  $z$  falls within the ellipsoid boundary  $E_j$  by verifying

$$\|\varphi_j(z)\|_2 \leq \Delta_j, \quad (18)$$

which implies  $z \in E_j$ . Otherwise, if

$$\|\varphi_j(z)\|_2 > \Delta_j, \quad (19)$$

$z$  is rejected as an open sample based on the high-confidence inference that  $z \notin \bigcup_{k=1}^K E_k$ .

## Experiments

### Experimental Settings

**Datasets** To validate the effectiveness of our method, we use three widely recognized datasets, including two intent datasets Banking (Casanueva et al. 2020) and OOS (Larson et al. 2019) and a question classification dataset StackOverflow (Xu et al. 2015). Details are provided in the appendix.

**Baselines** We compare our method with the following open-world classification approaches: OpenMax (Bendale and Boulton 2016), Softmax-MSP (Hendrycks and Gimpel 2018), DOC (Shu, Xu, and Liu 2017), ADB (Zhang, Xu, and Lin 2021), DA-ADB (Zhang et al. 2023a), KNNCL (Zhou, Liu, and Qiu 2022), DE (Zhou et al. 2023), CLAB (Liu et al.

2023) and MOGB (Li et al. 2025). Detailed descriptions of these baselines are provided in the appendix.

To ensure a fair comparison across different methods, we use BERT as the backbone model and conduct experiments on the open world classification platform TEXTTOIR (Zhang et al. 2021) for all methods. We integrated the released implementations of DE, CLAB and MOGB into TEXTTOIR. For other methods, the implementations included in TEXTTOIR are used.

**Evaluation Metrics** Following the metric adopted in previous work (Zhang, Xu, and Lin 2021), we use the average F1-score (F1) and the average accuracy (ACC) over all classes as metrics to evaluate the performance.

**Implementation Details** Following previous work DA-ADB (Zhang et al. 2023a), we randomly select 25%, 50%, and 75% of all classes as known classes with the remainder designated as unknown classes. The ratio of known classes selected is denoted as *Known Class Ratio* (KCR). Samples from unknown classes are excluded from the training and validation sets but retained in the test set, ensuring no unknown samples are present in the training and validation processes. For each setting, we conduct five rounds of experiments using random seeds 0,1,2,3 and 4 respectively and report the averages. Further implementation details of each method are provided in the appendix.

## Main Results

Results are presented in Table 1. Our method achieves the top results in almost all settings, demonstrating its effectiveness in enhancing both known and open intent classification. To assess generalization, we calculated the average performances across three KCR settings for each dataset. Our approach achieves 0.54%-2.66% improvement across all settings, highlighting the effectiveness and generalizability of the ellipsoid boundary and dual loss methodology in balancing empirical and open-space risks. Notably, compared to existing boundary-based methods ADB, DA-ADB and CLAB, our method exhibits performance gains specifically under 25% and 50% KCR settings. This improvement can be attributed to the ellipsoid boundary, which more effectively adapts to the presence of numerous unknown classes, resulting in a more precise decision boundary. This observation further suggests that our ellipsoid boundary design effectively improves the model’s detection performance.

## Discussions

To further validate the theory and demonstrate the effectiveness of our method, we conduct additional experiments and deep analysis from the following aspects: (1) the geometric advantage of ellipsoids comparing with balls, which is the primary novelty of our method; (2) ablation studies of major components; (3) the comparison with LLMs, showing the limitations of LLMs and supporting the value of our work. More discussions are stated in the appendix.

### Geometric Advantages of Ellipsoid

To show the geometric advantages of ellipsoids, we compare them against balls with different radii. To avoid bias in-

Datasets	Methods	KCR=25%		KCR=50%		KCR=75%		Average	
		F1	ACC	F1	ACC	F1	ACC	F1	ACC
Banking	Softmax-MSP	50.72	40.35	69.04	53.80	83.12	74.21	67.63	56.12
	DOC	62.09	61.19	80.43	75.30	<u>86.78</u>	81.23	76.43	72.57
	OpenMax	56.19	53.00	74.40	65.01	84.88	77.87	71.82	65.29
	ADB	73.21	81.50	81.09	79.35	85.61	80.72	79.97	80.52
	DA-ADB	74.75	81.06	82.22	80.76	85.69	80.89	<u>80.89</u>	<u>80.90</u>
	KNNCL	<u>77.30</u>	<u>85.72</u>	81.10	<u>82.61</u>	82.22	77.35	80.21	81.89
	DE	67.36	69.99	79.20	<u>72.64</u>	86.77	80.51	77.78	74.38
	CLAB	72.85	79.04	<u>83.17</u>	81.36	86.70	<u>81.92</u>	80.91	80.77
	MOGB	64.34	70.56	73.37	71.80	71.98	66.71	69.90	69.69
<i>EliDecide</i>	<b>77.75</b>	<b>85.81</b>	<b>83.74</b>	<b>82.90</b>	<b>86.80</b>	<b>81.97</b>	<b>82.76</b>	<b>83.56</b>	
OOS	Softmax-MSP	51.36	53.05	77.99	74.30	84.29	76.91	71.21	68.09
	DOC	72.32	81.32	84.51	84.18	<u>89.56</u>	86.90	82.13	84.13
	OpenMax	64.89	72.33	80.03	80.99	<u>74.58</u>	77.06	73.17	76.79
	ADB	79.18	89.03	85.77	87.21	88.83	86.85	84.59	87.70
	DA-ADB	80.51	89.85	<u>86.62</u>	88.84	88.53	87.47	<u>85.22</u>	88.72
	KNNCL	<u>82.05</u>	<b>92.95</b>	<u>84.51</u>	<u>89.07</u>	85.03	84.96	<u>83.86</u>	<u>88.99</u>
	DE	72.54	80.85	83.51	82.91	88.93	85.67	81.66	83.14
	CLAB	78.24	87.25	86.40	88.44	89.42	<u>88.02</u>	84.69	87.90
	MOGB	72.34	82.15	75.31	77.86	71.52	<u>71.47</u>	73.06	77.16
<i>EliDecide</i>	<b>84.48</b>	<u>92.57</u>	<b>88.57</b>	<b>90.56</b>	<b>90.21</b>	<b>88.98</b>	<b>87.75</b>	<b>90.70</b>	
StackOverflow	Softmax-MSP	38.13	27.69	62.30	49.50	76.17	69.99	58.87	49.06
	DOC	45.84	38.88	64.88	53.81	78.14	72.01	62.95	54.90
	OpenMax	45.37	38.49	69.68	63.44	79.82	74.62	64.96	58.85
	ADB	80.19	87.19	84.72	85.60	86.10	82.98	83.67	85.26
	DA-ADB	<u>83.44</u>	<u>90.02</u>	<u>86.77</u>	<u>87.82</u>	87.14	84.05	<u>85.78</u>	<u>87.30</u>
	KNNCL	77.47	83.26	84.70	84.84	87.10	83.87	83.09	83.99
	DE	60.86	60.12	75.92	71.58	84.18	79.78	73.65	70.49
	CLAB	73.61	78.07	85.79	86.19	<u>87.64</u>	<u>84.48</u>	82.35	82.91
	MOGB	70.91	76.78	79.13	77.91	82.47	78.06	77.50	77.58
<i>EliDecide</i>	<b>84.35</b>	<b>90.74</b>	<b>87.06</b>	<b>87.86</b>	<b>88.00</b>	<b>84.93</b>	<b>86.47</b>	<b>87.84</b>	

Table 1: Results of the main experiments with different Known Class Ratio (KCR) settings on three datasets. To compare the generalization of methods, the average performances across three KCR settings are also reported in the last two columns for each dataset. The best performances are shown in **bold**, while the second-best are underlined.

roduced by potentially suboptimal training procedures, we vary the radii by a hyperparameter called the *Coverage Fraction* (CF) instead of determining them through training. For each known class  $k$ , the ball’s radius is set as the distance from its centroid to the  $\lfloor \text{CF} |S_k| \rfloor$ -th nearest training sample. Consequently, each ball boundary is designed to include a fraction CF of its class’s training samples.

Results are presented in the top section of Table 2. Ball boundaries defined by CF fail to achieve optimal performance across different KCRs under the same hyperparameter, which highlights the necessity for adaptive training methods. Figure 4 illustrates the trade-off that a small CF causes excessive rejection of known samples as open, while a large CF fails to reject many open samples. Notably, even with KCR-specific CF tuning, balls underperform ellipsoids by 0.8%-1.9%. We attribute this gap to the inherent geometric constraint of balls, underscoring the advantage of the more flexible ellipsoidal shape.

## Ablation Studies

In this section we validate the effectiveness of our representation learning method and the negative loss function. The

results on OOS are shown in Table 2.

**Representation Learning** To explore the effect of our representation learning method, we discuss the effects of supervised contrastive learning (SCL) and normalization. Replacing SCL with a linear classifier and Softmax loss function reduces performance by 0.46%-1.33%, as the Softmax loss function solely focuses on distinguishing between different known classes but lacks intra-class compactness and inter-class separation. Removing feature normalization causes a 1.56%-2.89% drop at 25% and 50% known class ratios while flat at 75%. These results suggest discriminative gains are stronger when known samples are scarce, likely because compact known-class distributions are crucial with limited information and abundant unknowns.

**Negative Loss Functions** To validate the effectiveness of our proposed negative loss Eq. 15 which leverages generated pseudo-open samples, we compare it with the existing spherical boundary loss  $\mathcal{L}_{\text{ADB}}$  (Zhang, Xu, and Lin 2021) and two variants: CLAB loss (Liu et al. 2023)  $\mathcal{L}_{\text{CLAB}}$  and a variant ADB loss  $\mathcal{L}_{\text{ADB-Gen}}$  employing pseudo-open samples.  $\mathcal{L}_{\text{ADB}}$  employs known samples inside the boundary as

Open/Known	626	514	870	270	969	171	940	200
	51	4509	159	4401	356	4204	198	4362
	Known Ball (CF=0.8)	Open Ball (CF=0.8)	Known Ball (CF=0.975)	Open Ball (CF=0.975)	Known Ball (CF=1)	Open Ball (CF=1)	Known Ellipsoid	Open Ellipsoid

Figure 4: Coarse-grained confusion matrices with balls defined by different coverage fractions (CFs) and the ellipsoid (KCR=25%, seed=0).

Methods	25%	50%	75%
Ball (CF=0.8)	72.29	71.62	72.70
Ball (CF=0.9)	79.25	79.13	80.24
Ball (CF=0.95)	82.24	83.17	84.89
Ball (CF=0.975)	<u>83.20</u>	85.05	87.04
Ball (CF=0.9875)	82.49	86.00	88.25
Ball (CF=1)	80.07	<u>86.64</u>	<u>89.28</u>
Repre.: Softmax (w/o SCL)	83.78	87.25	89.67
Repre.: w/o Normalization	81.46	87.02	<b>90.80</b>
Loss: $\mathcal{L}_{ADB}$	37.42	59.21	74.19
Loss: $\mathcal{L}_{CLAB}$	42.49	60.67	74.53
Loss: $\mathcal{L}_{ADB-Gen}$	54.15	83.71	86.52
<i>EliDecide</i>	<b>84.48</b>	<b>88.57</b>	<u>90.21</u>

Table 2: Additional experimental results (F1) on OOS.

negative samples and its variant  $\mathcal{L}_{CLAB}$  extends it by using extra known samples from other classes as negative samples.  $\mathcal{L}_{ADB-Gen}$  replaces known samples with generated pseudo open samples inside the boundary as negatives, which is equivalent to setting  $\beta = 0$  in Eq. 15 to nullify the exponential term. Formal definitions are provided in the appendix.

Performance decreases sharply by 3%-47% particularly under low KCR settings. We attribute this reduction to  $\mathcal{L}_{ADB}$ 's reliance solely on intra-class information of known samples in a single class, which neglects critical inter-class relationships. While  $\mathcal{L}_{CLAB}$  mitigates this by using samples from other classes, improvements remain marginal.  $\mathcal{L}_{ADB-Gen}$  further incorporates the generated pseudo-open samples to model inter-class relationships but only activates those located inside the boundaries. In contrast, our method leverages the information across known classes to generate pseudo-open samples and activates all samples through the exponential term, better modeling the distribution of true open samples and leading to improved performance.

### Comparison with LLMs

While LLMs have demonstrated strong capabilities across various natural language understanding tasks, the capabilities on open world classification has not been sufficiently investigated. In this section, we evaluate Llama 3 (AI@Meta 2024) on open intent classification using multiple approaches and compare the results with our proposed method, highlighting the advantages of our approach using the relatively small-scaled pre-trained language model.

Experiments are conducted on Banking (Casanueva et al. 2020) with KCR=25%. We evaluate combinations of dif-

Detection Methods	Base		Finetuned	
	F1	ACC	F1	ACC
End-to-End(zero-shot)	38.03	23.60	45.13	26.00
End-to-End(five-shot)	40.60	29.60	44.86	25.41
Boundary	24.84	54.24	73.76	83.70

Table 3: Results of the experiments on Llama 3 with different backbones and detection methods on Banking with KCR=25%, where our proposed method's performance is F1=77.75 and ACC=85.81.

ferent backbone models and detection methods. Regarding backbone models, we employ both the pre-trained model without modification and a finetuned variant adapted to training data for enhanced discriminability on known classes. For the detection methods, we implement two distinct approaches. First, we utilize the LLM as a direct generator through tailored prompts, requiring the model to either select one of  $K$  known classes or classify the input as unknown (i.e. a  $(K + 1)$ -way classification problem). This approach is evaluated under both zero-shot and five-shot settings. Second, we employ the LLM as a feature extractor analogous to BERT. The final token embeddings serve as sentence-level representations ( $x_i$  in Eq. 1), with subsequent predictions made by the ellipsoid boundaries. More implementation details are stated in the appendix.

As shown in Table 3, finetuning on the known-class data improves classification performance. However, compared to specialized open intent detection methods, the end-to-end detection capability remains limited due to the absence of open samples during finetuning, consistent with the harmful *finetuning-induced degradation of vocabulary* phenomenon revealed in Zhou, Hong, and Qiu (2023). Notably, when substituting BERT with Llama 3 as the backbone, the results drop despite its larger scale. We attribute this to *overthinking* in deep transformer layers (Zhou et al. 2023), which may impair open-world discriminability. These results show the unique advantages of our proposed method using a compact encoder and a designed detection method.

### Conclusion

In this paper, we propose a novel ellipsoidal decision boundary framework EliDecide for open world classification. For each known class, we learn a nonsingular matrix that adaptively encodes both the ellipsoid's orientation and scale of its axes and optimize the boundaries via dual expansion and contraction losses. The geometric flexibility enables our method to jointly optimize empirical risk and open-space risk, achieving state-of-the-art performance across diverse benchmarks. Additional experiments validate each component, especially revealing the geometric advantage of ellipsoids over balls and the efficacy of our adaptive boundary methodology with a compact backbone encoder versus LLMs. Our results demonstrate that the ellipsoid's capacity to balance compactness and adaptability is critical for reconciling the dual challenges of preserving known-class fidelity while robustly rejecting unknown samples.

## Acknowledgments

This paper is funded by National Natural Science Foundation of China (Grant No. 62173195).

## References

- AI@Meta. 2024. Llama 3 Model Card.
- Bendale, A.; and Boulton, T. E. 2016. Towards open set deep networks. In *Proceedings of the IEEE conference on computer vision and pattern recognition*, 1563–1572.
- Casanueva, I.; Temčinas, T.; Gerz, D.; Henderson, M.; and Vulić, I. 2020. Efficient Intent Detection with Dual Sentence Encoders. In *Proceedings of the 2nd Workshop on Natural Language Processing for Conversational AI*, 38–45. Association for Computational Linguistics.
- Chen, G.; Xu, Q.; Zhan, C.; Wang, F. L.; Liu, K.; Liu, H.; and Hao, T. 2024. Improving Open Intent Detection via Triplet-Contrastive Learning and Adaptive Boundary. *IEEE Transactions on Consumer Electronics*, (1): 2806–2816.
- Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, 4171–4186. Association for Computational Linguistics.
- Gao, T.; Yao, X.; and Chen, D. 2021. SimCSE: Simple Contrastive Learning of Sentence Embeddings. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, 6894–6910.
- Gautam, C.; Parameswaran, S.; Kane, A.; Fang, Y.; Ramasamy, S.; Sundaram, S.; Sahu, S. K.; and Li, X. 2024. Class Name Guided Out-of-Scope Intent Classification. In *Findings of the Association for Computational Linguistics: EMNLP 2024*, 9100–9112. Association for Computational Linguistics.
- Geiger, A.; Lenz, P.; and Urtasun, R. 2012. Are we ready for autonomous driving? The KITTI vision benchmark suite. In *2012 IEEE Conference on Computer Vision and Pattern Recognition*, 3354–3361.
- Hendrycks, D.; and Gimpel, K. 2018. A Baseline for Detecting Misclassified and Out-of-Distribution Examples in Neural Networks. arXiv:1610.02136.
- Khosla, P.; Teterwak, P.; Wang, C.; Sarna, A.; Tian, Y.; Isola, P.; Maschinot, A.; Liu, C.; and Krishnan, D. 2020. Supervised Contrastive Learning. In *Advances in Neural Information Processing Systems*, 18661–18673. Curran Associates, Inc.
- Larson, S.; Mahendran, A.; Peper, J. J.; Clarke, C.; Lee, A.; Hill, P.; Kummerfeld, J. K.; Leach, K.; Laurenzano, M. A.; Tang, L.; and Mars, J. 2019. An Evaluation Dataset for Intent Classification and Out-of-Scope Prediction. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, 1311–1316. Association for Computational Linguistics.
- Li, Y.; Ouyang, X.; Pan, C.; Zhang, J.; Zhao, S.; Xia, S.; Yang, X.; Wang, G.; and Li, T. 2025. Multi-Granularity Open Intent Classification via Adaptive Granular-Ball Decision Boundary. *Proceedings of the AAAI Conference on Artificial Intelligence*, 39(23): 24512–24520.
- Liang, S.; Li, Y.; and Srikant, R. 2020. Enhancing The Reliability of Out-of-distribution Image Detection in Neural Networks. arXiv:1706.02690.
- Lin, T.-E.; and Xu, H. 2019. Deep Unknown Intent Detection with Margin Loss. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, 5491–5496. Association for Computational Linguistics.
- Liu, X.; Li, J.; Mu, J.; Yang, M.; Xu, R.; and Wang, B. 2023. Effective Open Intent Classification with K-center Contrastive Learning and Adjustable Decision Boundary. *Proceedings of the AAAI Conference on Artificial Intelligence*, (11): 13291–13299.
- Scheirer, W. J.; de Rezende Rocha, A.; Sapkota, A.; and Boulton, T. E. 2012. Toward open set recognition. *IEEE transactions on pattern analysis and machine intelligence*, (7): 1757–1772.
- Shu, L.; Xu, H.; and Liu, B. 2017. DOC: Deep Open Classification of Text Documents. In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, 2911–2916. Association for Computational Linguistics.
- Xu, J.; Wang, P.; Tian, G.; Xu, B.; Zhao, J.; Wang, F.; and Hao, H. 2015. Short text clustering via convolutional neural networks. In *Proceedings of the 1st Workshop on Vector Space Modeling for Natural Language Processing*, 62–69.
- Yang, C.; Cao, B.; Tao, C.; Yan, R.; and Wu, L. 2024a. JOSC: A Joint Model for Detecting Out-of-distribution Services based on Supervised Contrastive Learning. In *2024 IEEE International Conference on Web Services (ICWS)*, 728–736.
- Yang, J.; Zhou, K.; Li, Y.; and Liu, Z. 2024b. Generalized out-of-distribution detection: A survey. *International Journal of Computer Vision*, (12): 5635–5662.
- Zhang, H.; Li, X.; Xu, H.; Zhang, P.; Zhao, K.; and Gao, K. 2021. TEXTOIR: An Integrated and Visualized Platform for Text Open Intent Recognition. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing: System Demonstrations*, 167–174. Association for Computational Linguistics.
- Zhang, H.; Xu, H.; and Lin, T.-E. 2021. Deep open intent classification with adaptive decision boundary. In *Proceedings of the AAAI Conference on Artificial Intelligence*, 14374–14382.
- Zhang, H.; Xu, H.; Zhao, S.; and Zhou, Q. 2023a. Learning discriminative representations and decision boundaries for open intent detection. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 1611–1623.
- Zhang, H.; Zhou, Q.; Xu, H.; Su, J.; Evans, R.; and Gao, K. 2024. Multimodal Classification and Out-of-distribution Detection for Multimodal Intent Understanding. arXiv:2412.12453.

Zhang, S.; Bai, J.; Li, T.; Yan, Z.; and Li, Z. 2023b. Modeling Intra-class and Inter-class Constraints for Out-of-Domain Detection. In Wang, X.; Sapino, M. L.; Han, W.-S.; El Abbadi, A.; Dobbie, G.; Feng, Z.; Shao, Y.; and Yin, H., eds., *Database Systems for Advanced Applications*, 142–158. Cham: Springer Nature Switzerland. ISBN 978-3-031-30678-5.

Zhou, Y.; Hong, J.; and Qiu, X. 2023. Towards Open Environment Intent Prediction. In Rogers, A.; Boyd-Graber, J.; and Okazaki, N., eds., *Findings of the Association for Computational Linguistics: ACL 2023*, 2226–2240. Toronto, Canada: Association for Computational Linguistics.

Zhou, Y.; Liu, P.; and Qiu, X. 2022. KNN-contrastive learning for out-of-domain intent classification. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 5129–5141.

Zhou, Y.; Yang, J.; Wang, P.; and Qiu, X. 2023. Two birds one stone: Dynamic ensemble for ood intent classification. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, 10659–10673.

Zimmerer, D.; Full, P. M.; Isensee, F.; Jäger, P.; Adler, T.; Petersen, J.; Köhler, G.; Ross, T.; Reinke, A.; Kascenas, A.; Jensen, B. S.; O’Neil, A. Q.; Tan, J.; Hou, B.; Batten, J.; Qiu, H.; Kainz, B.; Shvetsova, N.; Fedulova, I.; Dylov, D. V.; Yu, B.; Zhai, J.; Hu, J.; Si, R.; Zhou, S.; Wang, S.; Li, X.; Chen, X.; Zhao, Y.; Marimont, S. N.; Tarroni, G.; Saase, V.; Maier-Hein, L.; and Maier-Hein, K. 2022. MOOD 2020: A Public Benchmark for Out-of-Distribution Detection and Localization on Medical Images. *IEEE Transactions on Medical Imaging*, 41(10): 2728–2738.