Enhancing Dialogue Systems with Adaptive Decision Boundaries and Multi-task Learning for Open Intent Recognition

Ms. Shruti Devlekar*, Mr. Om Gaydhane[1], Ms. Janhavi Khanvilkar[2], Mr. Kshitij Shidore[3], Mr. Amit Singh[4]
[*,1,2,3,4] *AI-DS Vivekanand Education Society's Institute of Technology (VESIT) Mumbai, India*

Abstract—Modern dialogue systems have made significant strides in understanding and responding to user queries. Despite these advancements, they still encounter challenges when dealing with diverse user intents, especially those not covered in their training data. This limitation often leads to inaccurate responses and misinterpretations, impeding the potential for meaningful interaction. Addressing this issue, our project introduces an innovative approach by integrating two sophisticated techniques: Adaptive Decision Boundary (ADB) for open intent detection and Multi-task Pre-training and Contrastive Learning with Nearest Neighbors (MTP-CLNN) for open intent discovery. ADB's key feature is its ability to dynamically adjust decision boundaries surrounding known intent clusters, thereby effectively identifying open intents that fall outside these clusters. To further enhance the accuracy of the model, we have upgraded the ADB model to Adaptive Decision Boundary Learning via Expanding and Shrinking (ADBES). This enhancement includes an update to the model's loss function, by introducing the concept of shrinking boundaries. This modification allows for a more precise encapsulation of known intents and a better differentiation from emerging or unknown ones. Through the combination of ADBES and MTP-CLNN, our pipeline not only accurately identifies known intents but also uncovers new intent categories, facilitating a more robust and adaptable dialogue system capable of evolving with user needs.

Keywords—Intent Recognition, Open intent detection, Open intent discovery, Banking77, ADB, MTP-CLNN.

I. INTRODUCTION

In the specialized context of banking sector dialogue systems, the accurate interpretation of user intents stands as a cornerstone for delivering personalized and efficient customer service. The task is relatively straightforward when dealing with common banking inquiries, such as requests for account balance checks or loan interest rates. However, the scenario becomes significantly more complex with the introduction of open intents, which encompass a wide range of user queries that do not fit neatly into predefined categories. These can range from nuanced questions about the intricacies of mortgage refinancing to specific concerns about fraud prevention measures or the details of conducting international transactions.

Open Intent Recognition (OIR) is segmented into two distinct components: detection and discovery. The initial component, open intent detection, is tasked with discerning established intent categories and pinpointing the presence of any non-standard, open intents. While it excels at identifying predefined intent classes, it does not extend to discovering specific new intent categories. Conversely, the open intent discovery component takes this a step further by categorizing these detected open intents into finely detailed clusters, albeit without the capacity to recognize established intent categories.

Figure 1 illustrates that while it is straightforward to identify specific user intents such as Top up limits and Card Linking, the variety and uncertainty of user needs suggest that pre-defined categories might not cover all scenarios, underscoring the value of distinguishing open intents from known ones to boost service quality. Despite significant advancements made using sophisticated methodologies across various benchmark datasets, challenges persist that impede further progress in Open Intent Recognition research such as the inability of both components to simultaneously recognize known intents and discover new, open ones, leaving OIR to largely remain a theoretical concept. This highlights the need for a framework that integrates both the detection and discovery modules, thereby streamlining the OIR process and enhancing its practical application.

User Utterances	Intent Label	Discovered Label
Do top-up limits exist?	Top up limits	-
Will my transfer be done soon?	Pending Transfer	-
How do I re-add a card to the app?	Card Linking	-
What is the exchange rate?	Exchange Rate	-
..	..	
How do I start trading Bitcoin on your platform?	**Open Intent**	Cryptocurrency Trading
Can I login using fingerprint authentication?	**Open Intent**	Biometric Authentication

Fig. 1. An example for Open Intent Recognition.

We introduce an Open Intent Recognition pipeline to identify known intents and discover new, open intents within interactive systems. This pipeline commences with dataset organization, tailored for both intent detection and discovery, followed by the setting of hyper-parameters and the establishment of a feature extraction framework. Once configured, the system is trained on labeled data to distinguish between known and open intents. The detection model isolates known intents, which then inform the training of the discovery model. This model proceeds to sort the open intents into specific clusters. Integrating the findings from both modules, the approach finalizes with the extraction of the top three defining keywords for each open-intent cluster, providing clear and actionable labels for each identified intent category.

II. LITERATURE

A. Open Intent Detection

Open intent detection, which seeks to accurately categorize known intents while simultaneously identifying novel, unknown intents, has garnered considerable interest for its utility in enhancing dialogue systems and question-answering platforms. This evolving field has transitioned from initial machine learning approaches, such as SVM, which relied on intensive feature engineering (Scholkopf et al., 2001; Tax and Duin, 2004; Rifkin and Klautau, 2004; Jain, Scheirer, and Boult, 2014), to more advanced deep learning techniques. These techniques, capable of autonomously learning deep semantic features from text, have set new benchmarks in the domain of open intent classification. Notable methodologies include the calibration of softmax output confidence by Brychcín and Kral (2017) for unknown intent detection, the use of adversarial learning by Yu et al. (2017) for generating training samples, and the application of generative adversarial networks by Ryu et al. (2018) to create and reject out-of-distribution samples.

In response to the challenges of predefining confidence scores for out-of-scope utterances, Lin and Xu (2019) introduced DeepUNK, employing a margin-based method for classifier training and leveraging the local outlier factor for unknown intent detection. This was followed by Zhang, Xu, and Lin (2021), who proposed a spherical decision boundary post-processing method with a pre-trained intent encoder. Prem et al. (2021) optimized a deep neural network-based intent classifier for unknown intent detection through multi-objective optimization, while Zhan et al. (2021) developed a discriminative classifier trained on pseudo outliers generated via self-supervision.

Parallelly, methods for open intent detection have been bifurcated into threshold-based and geometrical feature-based categories. Threshold-based techniques like MSP (Hendrycks and Gimpel, 2017), DOC (Shu et al., 2017), and OpenMax (Bendale and Boult, 2016) rely on probability thresholds to identify low-confidence samples as open intents after initial training on known intent classification tasks. Meanwhile, geometrical feature-based methods, such as DeepUnk (Lin and Xu, 2019) and ADB (Zhang et al., 2021a), utilize metric learning and boundary loss, respectively, to detect open intents as anomalies. These methods underscore a pivotal shift towards utilizing deep learning's capability for better feature discrimination and balanced representation learning in open intent classification, though challenges in fully exploiting these capabilities persist.

B. Open Intent Discovery

The journey into Open Intent Discovery (OID) has begun with foundational studies focusing on the application of unsupervised clustering for the extraction of features and the categorization of intents. Early innovators in the domain, such as Shi et al. (2018) with their use of auto-encoders for feature extraction, Perkins and Yang (2019) with their emphasis on analyzing the context of utterances, and Chatterjee and Sengupta (2020) with their efforts to enhance density-based models, have laid the groundwork for future exploration. Subsequent research by Haponchyk et al. (2018; 2021) delved into the capabilities and limitations of supervised clustering for intent classification, highlighting the challenges in accommodating emergent intent types. This scenario set the stage for the exploration of semi-supervised OID methods, as demonstrated by the contributions of Forman et al. (2015), Lin et al. (2020), and Zhang et al. (2021c), showcasing the advantage of integrating known with unidentified intents to unearth open intent classifications.

Progress in semi-supervised OID was furthered by Lin et al. (2020) through a strategy that incorporated supervised learning on recognizable intents via tasks related to sentence similarity, augmented by pseudo labeling for the categorization of unlabeled utterances to enhance embedding accuracy. In a similar vein, Zhang et al. (2021c) crafted a strategy that merged pre-training on identifiable intents with the application of k-means clustering to unlabeled utterances, inspired by the Deep Clustering framework (Caron et al., 2018), aiming to refine learning processes and ensure the alignment of clusters. This approach aligns with methodologies that initially sort utterances into identifiable and unknown categories to facilitate the identification of new intents among the latter, a step that significantly hinges on the precision of the preliminary classification as depicted by Vedula et al. (2020) and Zhang et al. (2021b).

The efficacy of multi-task pre-training and contrastive learning for clustering and enhancing representations, marking a departure from traditional methods reliant on extensive annotated data. Despite the success of large-scale pre-trained models, their application is challenged by linguistic discrepancies, spurring interest in dialogue-specific continual pre-training. Innovations by Zhang et al. (2020) and Vulic et al. (2021); Zhang et al. (2021e), which utilize task-related pre-training with a focus on intent detection, exemplify this shift. Drawing inspiration from Zhang et al. (2021d), this approach benefits OID by using publicly available intent datasets alongside domain-specific unlabeled data, aligning well with the nature of OID tasks that inherently include unlabeled utterances.

Moreover, the potential of contrastive learning in both computer vision and natural language domains is highlighted, showcasing its role in refining sentence embeddings through unsupervised techniques. Key contributions from Gao et al. (2021), Yan et al. (2021), and others underscore the value of contrastive loss in achieving an isotropic embedding space, with Kim et al. (2021) and Giorgi et al. (2021) further demonstrating its utility in improving BERT representations and developing universal sentence encoders. The effectiveness of self-supervised pre-training coupled with supervised fine-tuning for few-shot intent recognition, as well as the integration of contrastive loss with clustering goals for enhancing short text clustering, reaffirms the advantage of fostering semantic cohesion among utterances while addressing the pitfalls of false negatives typical in contrastive learning schemes.

III. METHODOLOGY

The interconnection between open intent detection and discovery modules is pivotal, yet there exists no comprehensive framework to seamlessly activate both modules for the dual purposes of identifying established intents and unearthing novel, open intents. Our framework initiates with preprocessing the initial dataset, subsequently channeling the labeled data of known intents into the open intent detection module for model training, as chosen by the user. Given the substantial volume of unlabeled data, which likely encompasses both known and open intents, the framework employs a robustly trained open intent detection model to categorize the unlabeled dataset. The outcome of this process includes both the recognized known intents and the newly identified open intents.

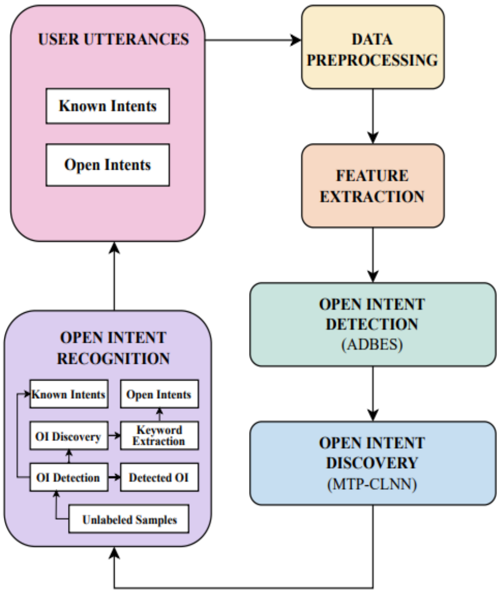

Fig. 2. Framework for Open Intent Recognition.

These results, alongside the original labeled dataset, are fed into the open intent discovery module, enhancing its training inputs through data augmentation.

User-selected clustering techniques are then applied within this module to delineate distinct clusters of open intents. Following the completion of training across both modules, the framework is adept at conducting open intent recognition on new, unlabeled datasets. The open intent detection module initially applies its refined model to segregate known intents from potential open intents within the data. Subsequently, the open intent discovery module steps in to categorize these open intents into more nuanced, fine-grained clusters. The framework extracts pivotal keywords from sentences within each open intent cluster, thereby assigning informative, keyword-based labels to each group. Through this sophisticated framework, we not only pinpoint established intents but also reveal and label new, open intents, enriching the dataset with valuable insights and classifications.

We implement a method for adaptively learning decision boundaries tailored for open intent classification through a dynamic process of expansion and contraction. This approach begins by determining a decision boundary for each recognized intent class, defining a central point (ck) and a variable radius (Δk) for the boundary around each intent class k. The center ck is calculated as the mean vector of all instance representations within class k, while the boundary's radius Δk is adjusted through learning to encompass as many relevant instances as possible, aiming to minimize both the empirical risk of excluding known intents and the open space risk of including too many.

To navigate the delicate equilibrium between minimizing these risks, we adopt the Adaptive Decision Boundary (ADB) technique, which fine-tunes the radius of the decision boundaries by balancing the need to expand the boundary to include as many known intent instances as possible against the need to contract it to exclude outliers. This balance is achieved through a loss function that adjusts the boundary radius based on the proximity of each instance to the class center, modulated by a variable γ that shifts emphasis between expanding and contracting the boundary based on whether an instance falls inside or outside the current radius.

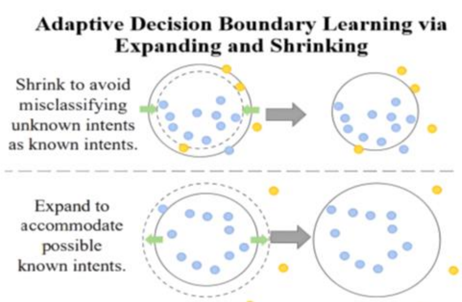

Fig. 3. Working of ADBES model.

To refine this process further, we implement a strategy of boundary adjustment by taking into account instances outside the current class (negative instances), allowing for a more nuanced determination of the decision boundary. This involves either expanding the boundary to include more instances that are closely aligned with the class intent or contracting it to exclude negative instances that are too close. By introducing parameters for expansion (e) and contraction (s), we dynamically adjust the radius based on the proximity of negative instances to the decision center, encapsulated in a loss function that incorporates these adjustments to fine-tune the decision boundaries.

This enhanced loss function involves distinguishing between instances falling outside (positive loss) and inside (negative loss) the decision boundary. Instances further away from the centroid than the boundary radius contribute to the positive loss, encouraging the expansion of the boundary to encompass these instances. Conversely, instances within the boundary radius contribute to the negative loss, which is scaled down for instances closer to the centroid, promoting a contraction of the boundary to exclude outliers.
Our loss function intricately merges positive and negative loss components to refine decision boundaries for open intent classification, achieving a delicate balance between the inclusivity of known intents and the exclusion of outliers. This adaptive mechanism tailors decision boundaries dynamically, considering the distribution of both in-class and out-of-class instances, to enhance the precision and effectiveness of classifying open intents. By embodying our expanding and shrinking methodology, this approach ensures decision boundaries are adjusted in a nuanced manner, informed by the proximity of instances to class centroids. Consequently, it fosters the development of decision boundaries that accurately encompass a majority of known intent instances while adeptly identifying and segregating open intent instances. This methodology underscores a vital equilibrium between embracing relevant instances for robust intent recognition and maintaining specificity by filtering out unrelated data, thereby optimizing the classification process for both known and emerging open intents.

IV. EXPERIMENTS

A. Dataset

We conduct experiments on three benchmark datasets: *BANKING*, with 13,083 queries across 77 intents (Casanueva et al., 2020); *OOS*, featuring 150 intents, 22,500 in-domain, and 1,200 out-of-domain queries (Larson et al., 2019); and a curated version of *StackOverflow*, which has 20 classes and a total of 20,000 samples from over 3.37 million questions (Xu et al., 2015).

Dataset	Class	Train/Valid/Test	Length (max/mean)
BANKING	77	9003 / 1000 / 3080	79 / 11.91
OOS	150	15000 / 3000 / 5700	28 / 8.31
StackOverflow	20	12000 / 2000 / 6000	41 / 9.18

TABLE I. Statistics of Datasets

B. Baselines

Our approach is compared against five methods in open intent classification. These include MSP (Hendrycks and Gimpel 2016), which utilizes softmax predictions for identifying out-of-distribution instances, and DOC (Shu, Xu, and Liu 2017), employing deep learning to craft a multi-class classification framework while enhancing decision boundaries through Gaussian fitting. Additionally, OpenMax (Bendale and Boult 2016) integrates meta-recognition with activation patterns in the penultimate layer to mitigate open space risks, whereas DeepUnk (Lin and Xu 2019) shifts from softmax to margin loss to refine deep feature discrimination, emphasizing the separation between classes and cohesion within them. Finally, ADB (Zhang, Xu, and Lin 2021) introduces an innovative post-processing technique for categorizing open intents by establishing a spherical decision boundary around each recognized class.

C. Evaluation Metrics

Building on the approach by Zhang, Xu, and Lin (2021), we employ accuracy and macro F1-score as the primary metrics to assess performance in open intent classification. All unclassified intents are aggregated into a single open class for evaluation purposes. The overall accuracy and macro F1-score across all intents are calculated and represented as Acc and F1, respectively. Furthermore, to thoroughly assess the classifiers' effectiveness in distinguishing between known and unknown intents, we separately calculate the macro F1-scores for both known intents (indicated as Known) and unknown intents (indicated as Unknown).

D. Implementation Details

For the ADBES model, In alignment with the methodologies outlined by Shu, Xu, and Liu (2017) and Lin and Xu (2019a), our approach involves designating a subset of classes as unknown (open) for the duration of the training phase, reintroducing them only during the testing phase. We partition all datasets into three distinct sets: training, validation, and testing. Within the training set, we vary the proportion of known classes to include 25%, 50%, and 75%, respectively, treating the remainder as a singular open class that is excluded from training. Both the identified known classes and the aggregated open class are then incorporated into the testing phase. We calculate the mean performance across ten experimental iterations for each specified ratio of known classes. For our model's architecture, we utilize the BERT (bert-uncased, featuring a 12-layer transformer) framework as implemented in PyTorch, as per Wolf et al. (2019), adhering closely to the recommended optimization hyperparameters. To enhance training efficiency and performance outcomes, we opt to freeze the parameters across all but the final transformer layer of BERT. The model is trained with a batch size of 128 and a learning rate of 2e-5. For optimizing the boundary loss (Lb), we employ the Adam optimization algorithm (Kingma and Ba, 2014) with a specific learning rate of 0.05 to fine-tune the boundary parameters.

For the MTP-CLNN model, We base our model on the bert-base-uncased pre-trained model provided by Wolf et al. (2019), utilizing the [CLS] token as the representation output by BERT. For Multi-Task Processing (MTP), the model is initially trained to convergence using an external dataset. In the domain of contrastive learning, we transform the 768-dimensional BERT embeddings into 128-dimensional vectors using a two-layer Multi-Layer Perceptron (MLP) and adjust the softmax temperature to 0.07. For the identification of nearest neighbors, we employ the inner product strategy as recommended by Johnson et al. (2017). The size of the neighborhood (K) is set to 50 for the BANKING dataset, based on the empirical observation that the ideal K value is approximately half the average training set size per class. This neighborhood configuration is refreshed every five epochs. Regarding data augmentation, we set a random token replacement probability at 0.25. The AdamW optimizer, again referenced from Wolf et al. (2019), is used for model optimization. In the initial stage, the learning rate is 5e-5, while in the second stage, it is adjusted to 1e-5 for the BANKING. All experiments were executed on a V100 GPU.

V. RESULTS

A. Visualization

Figure 4 visualizes the distribution of intent center clusters, distinguishing between known and open classes. These visualizations underscore the MTP-CLNN model's capacity to differentiate intents, which is pivotal for enhancing intent recognition accuracy.

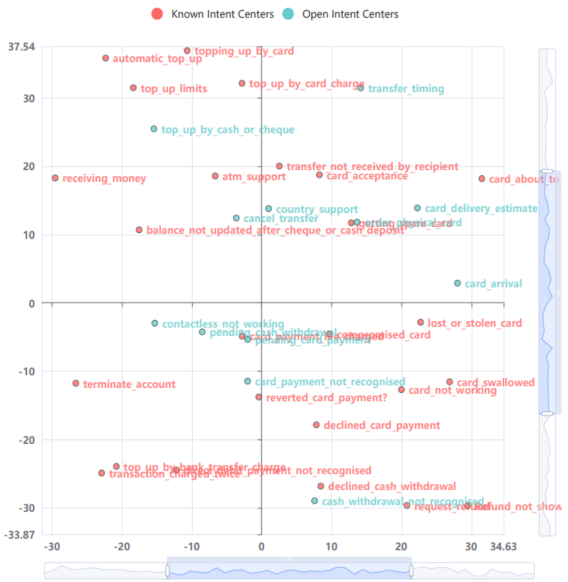

Fig. 4. Visualization of Intent center distribution.

The Open Intent Recognition pipeline is depicted in Figure 5 below. It processes a user query, identifying known intents. Queries classified as open intents are then directed to the MTP-CLNN model for open intent discovery.

```
Enter your query: Is Visa or Mastercard available?
Predicted category by ADB:  visa_or_mastercard
Enter your query: What is the auto top-up limit?
Predicted category by ADB:  automatic_top_up
Enter your query: Someone stole my phone.
Predicted category by ADB:  lost_or_stolen_phone
Enter your query: I want to play football with friends
Detected as Open Intent by ADB. Further analyzing with MTP-CLNN...
Predicted Intent by MTP-CLNN:  play_friends_want
Enter your query: exit
Exiting...
```

Fig. 5. Open Intent Recognition pipeline.

B. *Result Analysis*

Table II displays the macro F1 scores for Open Intent Detection, comparing overall intent classes ("ALL"), unknown (open) intent classes ("Unknown"), and known classes ("Known"). Our method notably surpasses baseline methods, demonstrating its effectiveness. Particularly, it outperforms the ADB method in the BANKING dataset by margins of 1.24%, 2.12%, and 1.38% for 25%, 50%, and 75% known class proportions, respectively. This improvement is observed not just for unknown classes but also for known classes, suggesting our method's capability to define precise and appropriate decision boundaries for different intent classes.

Known cls ratio	BANKING				
	Methods	All	Unknown	Known	Accuracy
25%	MSP	50.09	41.43	50.55	43.67
	DOC	58.03	61.42	57.85	56.99
	OpenMax	54.14	51.32	54.28	49.94
	DeepUnk	61.36	70.44	60.88	64.21
	ADB	71.62	84.56	70.94	78.85
	ADBES	**72.86**	**85.56**	**73.19**	**80.75**
50%	MSP	71.18	41.19	71.97	59.73
	DOC	73.12	55.14	73.59	64.81
	OpenMax	74.24	54.33	74.76	65.31
	DeepUnk	77.53	69.53	77.74	72.73
	ADB	80.9	78.44	80.96	78.86
	ADBES	**83.02**	**81.43**	**83.96**	**80.77**
75%	MSP	83.6	39.23	84.36	75.89
	DOC	83.34	50.6	83.92	76.77
	OpenMax	84.07	50.85	84.64	77.45
	DeepUnk	84.31	58.54	84.75	78.52
	ADB	85.96	66.47	86.29	81.08
	ADBES	**87.34**	**68.56**	**88.02**	**83.15**

TABLE II. The results of open intent classification by varying the proportions (25%, 50% and 75%) of known classes.

The ADBES method establishes more precise decision boundaries for open intent classification, showing significant effectiveness, particularly with a 25% known class proportion. Its capacity to discern open intents notably surpasses that of the ADB method.

The development of IntelliBank, a Service Bot designed for Open Intent Recognition in banking conversations, represents a significant application of our findings. By integrating Adaptive Decision Boundary Learning with Expanding and Shrinking for precise Open Intent Detection, alongside Multi-task Pre-training and Contrastive Learning with Nearest Neighbors for Open Intent Discovery, IntelliBank substantially improves intent recognition accuracy. It effectively discerns user queries and, upon detecting open intents, suggests relevant keywords. This facilitates efficient query resolution by bank staff, optimizing operations and enhancing customer service. The results indicate that IntelliBank's application of the ADBES method and MTP-CLNN model contributes significantly to its performance, underscoring the practical value of our research in real-world settings.

Notably, our model also demonstrates its efficacy on other datasets like OOS and StackOverflow. It successfully predicts known intents and identifies open intents, confirming its adaptability and effectiveness across different domains.

VI. CONCLUSION

Our research introduces a significant advancement in intent recognition technology, particularly in banking contexts, through the development of a sophisticated intent recognition pipeline. Central to our approach is the Adaptive Decision Boundary (ADB) method, enhanced by a novel adjustment to its loss function that incorporates expanding and shrinking boundaries to accurately exclude out-of-class examples. This refinement has led to a notable increase in accuracy of approximately 2% across 25%, 50%, and 75% known class ratios. The Multi-task Pre-training and Contrastive Learning with Nearest Neighbors (MTP-CLNN) model further complements this by utilizing a dual-stage strategy. Initially, it leverages both external and internal data for comprehensive representation learning. Subsequently, it employs contrastive learning to harness self-supervisory signals, significantly bolstering the performance of New Intent Detection (NID) under both unsupervised and semi-supervised conditions.

Moreover, our contribution includes the development of an open intent recognition pipeline that combines the ADBES and MTP-CLNN model’s strengths. This pipeline effectively identifies known intents and uncovers new ones, showcasing significant effectiveness on the BANKING77 dataset. Additionally, the model's versatility has been further validated on other datasets such as OOS and StackOverflow, highlighting its widespread applicability and strong performance in various contexts.

Future directions for this research include exploring the scalability of the ADBES method across various domains, enhancing the MTP-CLNN model for real-time processing in dynamic environments, and integrating advanced NLP technologies for automated response generation.

CONFLICT OF INTEREST

On behalf of all authors, the corresponding author states that there is no conflict of interest.

FUNDING INFORMATION

The authors received no specific funding for this work.

AUTHOR CONTRIBUTION

Shruti Devlekar and Om Gaydhane conceived the study and worked on the algorithms. Janhavi Khanvilkar designed the user interface. Mr. Amit Singh provided guidance for the research. All authors contributed equally to writing the manuscript.

DATA AVAILABILITY STATEMENT

All relevant data are within the paper.

RESEARCH INVOLVING HUMAN AND/OR ANIMALS

Not Applicable.

INFORMED CONSENT

Not Applicable.

REFERENCES

[1]. Zhang, H.; Xu, H.; Lin, T.-E.; and Lyu, R. 2021. Deep open intent classification with adaptive decision boundary. In Proceedings of the AAAI Conference on Artificial Intelligence, volume 35, 14374–14382.

[2]. Hendrycks, D.; and Gimpel, K. 2017. A baseline for detecting misclassified and out-of-distribution examples in neural networks. arXiv preprint arXiv:1610.02136.

[3]. Y. Zhang, H. Zhang, L.-M. Zhan, X.-M. Wu, and A. Y. S. Lam, "New Intent Discovery with Pre-training and Contrastive Learning", in Proceedings of the 2022 Annual Conference of the Association for Computational Linguistics (ACL 2022), Dublin, May 2022.

[4]. Casanueva, I.; Temcinas, T.; Gerz, D.; Henderson, M.; and Vulic, I. 2020. Efficient intent detection with dual sentence encoders. arXiv preprint arXiv:2003.04807.

[5]. X. Liu, J. Li, J. Mu, M. Yang, R. Xu, and B. Wang, "Effective Open Intent Classification with K-Center Contrastive Learning and Adjustable Decision Boundary", in Proceedings of The AAAI Conference on Artificial Intelligence (AAAI), 2023.

[6]. H. Zhang, X. Li, H. Xu, P. Zhang, K. Zhao, and K. Gao, "TEXTOIR: An Integrated and Visualized Platform for Text Open Intent Recognition", ACL 2021, Demo paper, pp. 167-174.

[7]. Devlin, J.; Chang, M.-W.; Lee, K.; and Toutanova, K. 2019. BERT: Pre-training of Deep Bidirectional Transformers for Language Understanding. ArXiv, abs/1810.04805.

[8]. Gao, T.; Yao, X.; and Chen, D. 2021. SimCSE: Simple Contrastive Learning of Sentence Embeddings. arXiv preprint arXiv:2104.08821.

[9]. Larson, S.; Mahendran, A.; Peper, J. J.; Clarke, C.; Lee, A.; Hill, P.; Kummerfeld, J. K.; Leach, K.; Laurenzano, M. A.; Tang, L.; et al. 2019. An evaluation dataset for intent classification and out-of-scope prediction. arXiv preprint arXiv:1909.02027.

[10]. Kang, B.; Li, Y.; Xie, S.; Yuan, Z.; and Feng, J. 2020. Exploring balanced feature spaces for representation learning. In the International Conference on Learning Representations.

[11]. Ryu, S.; Kim, S.; Choi, J.; Yu, H.; and Lee, G. G. 2017. Neural sentence embedding using only in-domain sentences for out-of-domain sentence detection in dialog systems. Pattern Recognition Letters, 88(Mar.1): 26–32.

[12]. Goodfellow, I.; Pouget-Abadie, J.; Mirza, M.; Xu, B.; Warde-Farley, D.; Ozair, S.; Courville, A.; and Bengio, Y. 2014. Generative adversarial nets. Advances in neural information processing systems, 27.

[13]. Guanhua Chen, Qiqi Xu, Choujun Zhan, Fu Lee Wang, Kai Liu, Hai Liu, Tianyong Hao. "Improving Open Intent Detection via Triplet-Contrastive Learning and Adaptive Boundary", IEEE Transactions on Consumer Electronics, 2024.

[14]. Hua Xu, Hanlei Zhang, Ting-En Lin. "Chapter 10 Experiment Platform for Dialogue Intent Recognition Based on Deep Learning", Springer Science and Business Media LLC, 2023.

[15]. Wenbin An, Feng Tian, Wenkai Shi, Haonan Lin, Yaqiang Wu, Mingxiang Cai, Luyan Wang, Hua Wen, Lei Yao, Ping Chen. "DOWN:Dynamic Order Weighted Network for Fine-grained Category Discovery", Knowledge-Based Systems, 2024.