

ProtoInfoMax: Prototypical Networks with Mutual Information Maximization for Out-of-Domain Detection

Iftitahu Ni'mah^{1,2}, Meng Fang¹, Vlado Menkovski¹, Mykola Pechenizkiy¹

Eindhoven University of Technology (TU/e), Eindhoven, The Netherlands¹

Research Center for Informatics (BRIN), Bandung, Indonesia^{2*}

{i.nimah, m.fang, v.menkovski, m.pechenizkiy}@tue.nl¹

Abstract

The ability to detect Out-of-Domain (OOD) inputs has been a critical requirement in many real-world NLP applications. For example, intent classification in dialogue systems. The reason is that the inclusion of unsupported OOD inputs may lead to catastrophic failure of systems. However, it remains an empirical question whether current methods can tackle such problems reliably in a realistic scenario where zero OOD training data is available. In this study, we propose ProtoInfoMax, a new architecture that extends Prototypical Networks to simultaneously process in-domain and OOD sentences via Mutual Information Maximization (InfoMax) objective. Experimental results show that our proposed method can substantially improve performance up to 20% for OOD detection in low resource settings of text classification. We also show that ProtoInfoMax is less prone to typical overconfidence errors of Neural Networks, leading to more reliable prediction results.¹

1 Introduction

Many real-world applications imply an open world assumption (Scheirer et al., 2013; Fei and Liu, 2016)², requiring intelligent systems to be aware of novel Out-of-Domain (OOD) examples, given limited In-Domain (ID) and zero OOD training data. Intent classification for conversational AI services, for instance, may have to deal with unseen OOD utterances (Tan et al., 2019; Kim and Kim, 2018a; Larson et al., 2019; Zheng et al., 2020). Question answering system is also preferred to have a certain degree of language understanding via its

*Komplek LIPI, Jl. Sangkuriang, Dago, Coblong, Bandung, Indonesia 40135. Phone: (+6222) 2504711.

¹Code and preprocessed data are available at <https://github.com/inimah/protoinfomax.git>.

²System built under this assumption should be able to not only correctly analyze In-Domain (ID) inputs but also reliably reject Out-of-Domain (OOD) inputs that are not supported by the system.

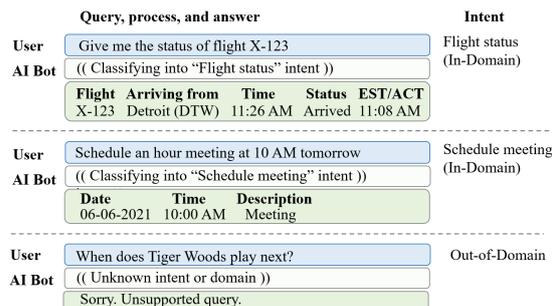


Figure 1: An example of OOD detection in task-oriented dialog systems.

ability to contrast between relevant and irrelevant sentences (Yeh and Chen, 2019). Likewise, a classifier trained on past topics of social media posts is often expected to be aware of future social media streams with new unseen topics (Fei and Liu, 2016; Fei et al., 2016). An example of an AI system with OOD awareness is illustrated in Figure 1. When a user inputs an unknown query with OOD intent, instead of providing random feedback, a system that is aware of OOD inputs can better respond informatively.

To develop a reliable intelligent system that can correctly process ID inputs and detect unclassified inputs from different distribution (OOD), existing approaches often formulate OOD detection as anomaly detection (Ryu et al., 2017, 2018; Hendrycks et al., 2019). The concept of learning ID classification and OOD detection tasks simultaneously is also incorporated in diverse applications, including open text classification (Shu et al., 2017) and OOD detection in task-oriented dialog system (Kim and Kim, 2018b; Zheng et al., 2020). These methods rely on large-scale ID and OOD labeled training data or well-defined data distributions.

Unfortunately, large data settings make the methods unrealistic for many real world applications with limited ID and zero OOD training data. As a result, current research introduces few-shot and zero-shot learning frameworks for OOD detection

problems in a low resource scenario of text classification (Tan et al., 2019). Their objective is to learn a metric space for ID and OOD prediction given prototype representation of ID sentences and target sentences sampled from ID and OOD distribution. However, the current method neglects an overconfidence issue of the trained Prototypical Networks in the inference stage where both novel ID and OOD inputs occur. For example, OOD samples are likely to be classified as ID with a high similarity score (Liang et al., 2018; Shafaei et al., 2019), especially if they share common patterns or semantics with ID samples (e.g. common phrases, sentence topicality, sentiment polarity) (Lewis and Fan, 2019).

To mitigate the above problems, we adopt Mutual Information Maximization (InfoMax) objective (Belghazi et al., 2018; Hjelm et al., 2019) for regularizing Prototypical Networks (Section 4.1). We extend Prototypical Networks (Snell et al., 2017) to learn multiple prototype representations by maximizing Mutual Information (MI) estimates between sentences that share a relevant context, such as keywords (Section 4.3). We demonstrate that our proposed method is less prone to typical overconfidence error of Neural Networks (Lakshminarayanan et al., 2017; Guo et al., 2017; Liang et al., 2018; Shafaei et al., 2019). This result leads to more reliable prediction outcomes, specifically in the inference stage where the model has to deal with both novel ID and OOD examples. Overall, experimental results on real-world low-resource sentiment and intent classification (Section 6) show that the proposed method can substantially improve performance of the existing approach up to 20%.

To summarize, our contributions are as follows:

- We introduce ProtoInfoMax – Prototypical Networks that learn to distinguish between ID and OOD representations via Mutual Information Maximization (InfoMax) objective (Section 4.1).
- We enhance ProtoInfoMax by incorporating multiple prototype representations (Section 4.3) to further improve the discriminability of the learned metric space.
- We further investigate the reliability of Prototypical Networks in this study, in addition to common metrics used for evaluating Out-of-Domain detection (Section 6.3-6.4). Our problem of interest is determining whether

the learned metric space indicates a well calibrated model. That is, a condition where the trained model assigns high similarity score for test samples drawn from ID distribution and assigns lower similarity score for samples drawn from OOD distribution.

2 Related Work

2.1 Few-shot Learning

Few-shot Learning (FSL) has been increasingly studied in NLP. Several works have adopted the experimental protocol of FSL, expanding the application of FSL in text classification (Yu et al., 2018; Bao et al., 2020; Tan et al., 2019; Zhang et al., 2020) and other tasks (Fang et al., 2017; Han et al., 2018; Gao et al., 2019; Sun et al., 2019; Dopierre et al., 2021).

2.2 Out-of-Domain Detection

The problem of Out-of-Domain (OOD) detection has been investigated in many contexts; such as anomaly detection (Zenati et al., 2018; Hendrycks et al., 2019), one-class classification (Khan and Madden, 2014; Ruff et al., 2018), open-set recognition (Geng et al., 2020), and novelty detection (Perera et al., 2019). In speech recognition and language understanding domain, the problem is formulated as OOD utterances detection (Lane et al., 2006; Tur et al., 2014; Zheng et al., 2020). Most of these works, however, depend on the availability of large-scale ID and OOD samples, in addition to the inclusion of OOD samples in training data as supervision signals for the model.

2.3 Mutual Information Objective

Incorporating Maximization Mutual Information (MMI) (Linsker, 1988; Bell and Sejnowski, 1995), which we refer to as InfoMax (section 4.1), as training objective for Neural Networks is exemplified by early work on diversifying neural conversational model (Li et al., 2016). However, the work mainly uses the MMI objective in the inference stage for controlling the decoder outputs. Recent works on InfoMax objective for deep learning (Belghazi et al., 2018; Hjelm et al., 2019) have introduced simple yet effective loss function approximation, such that the objective can be used in the training stage. Prior to our study, InfoMax objective has been adapted to learn useful representations by maximizing relevant information between local and global features of image data (Hjelm et al.,

2019), to learn speaker representations (Ravaneli and Bengio, 2019), and to learn robust question answering system (Yeh and Chen, 2019).

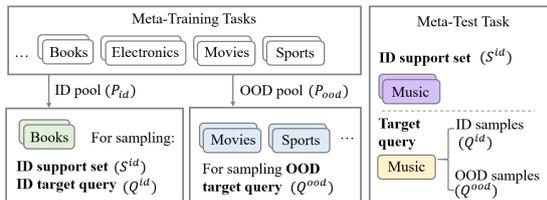


Figure 2: Illustration on how ID and OOD sentences are sampled during meta-training and meta-test tasks. (“Books”, “Electronics”, ...) exemplifies domains available in training data, disjoint with examples in test data.

3 Problem Definition

Similar to the previous setting (Tan et al., 2019), we consider the zero-shot OOD detection problem for meta-tasks in this study. In general, there are three main inputs for prototypical learning in this study: ID support set S^{id} , ID target query Q^{id} , and OOD target query Q^{ood} (Figure 2 and Figure 3a).

Meta-training tasks For each training episode, ID examples \mathcal{D}_{id} are drawn independently from ID distribution in meta-training tasks $\mathcal{P}_{T_i}, T_i \in \mathcal{T}$ (later refer to as \mathcal{P}_{id}). Specifically:

$$\mathcal{D}_{id} = \{(x_1, y_1), \dots, (x_n, y_n)\} \sim \mathcal{P}_{T_i}, \quad (1)$$

where each ID example is composed of sentences x and their corresponding ID classes y . In sentiment classification benchmark, this y is a sentiment label, $y \in \{\text{“positive” or “negative”}\}$. In Figure 2, \mathcal{P}_{T_i} or \mathcal{P}_{id} is described as ID domain “Books”. ID support set S^{id} and ID target query Q^{id} are drawn from \mathcal{D}_{id} , where S^{id} and Q^{id} are mutually exclusive: $S^{id} \cap Q^{id} = \emptyset$.

OOD data \mathcal{D}_{ood} is drawn from out-of-episode distribution \mathcal{P}_{T_j} :

$$\mathcal{D}_{ood} = \{(x_1, y_1), \dots, (x_n, y_n)\} \sim \mathcal{P}_{T_j}, \quad (2)$$

where tasks or domains T_j are disjoint with those in training: $T_j \in \mathcal{T}, T_j \neq T_i$ (later refer to as \mathcal{P}_{ood}). \mathcal{P}_{T_j} or \mathcal{P}_{ood} is described as out-of-scope domains (“Movies”, “Sports”, ...) in Figure 2.

Meta-test tasks As illustrated in Figure 2, each task or domain in the test set (e.g. “Music”) is composed of a disjoint ID support set and target queries (ID and OOD). ID examples are accompanied with ID class labels, i.e. $y \in \{\text{“positive”, “negative”}\}$ for

sentiment classification benchmark, while OOD examples are annotated with OOD labels ($y = \text{“ood”}$). Note that this $y = \text{“ood”}$ is unknown to the model during training.

4 ProtoInfoMax

We propose two models: **ProtoInfoMax** and **ProtoInfoMax++**, briefly illustrated in Figure 3. The main difference between the two models is their prototype generator, further discussed in Section 4.3. **ProtoInfoMax++** merges multiple prototype representations: (i) standard feature averaging prototype vector based on sentence features, referred to as C_S ; and (ii) prototype vector based on keyword context features, referred to as C_w . We regularize both models with an InfoMax objective, discussed in Section 4.1.

4.1 InfoMax Objective

We adopt the recently proposed Mutual Information Maximization (InfoMax) training objective for deep learning (Belghazi et al., 2018; Hjelm et al., 2019) as a contrastive view of data drawn from ID and OOD distribution. The idea is simple: we want to maximize Mutual Information (MI) estimates for samples drawn from ID distribution \mathcal{P}_{id} , while penalizing OOD samples with lower MI estimates.

Here, a multi-objective for simultaneously classifying ID sentences and detecting OOD sentences is formulated as a contrastive learning framework via an InfoMax objective. The model is enforced to learn binary reject function \mathcal{L} that partitions the input space \mathcal{X} with respect to \mathcal{P}_{id} and \mathcal{P}_{ood} . Incorporating a binary reject function for regularizing Prototypical Networks in the current OOD detection problem can simplify the overall training mechanism. Namely, it can be approximated by a simple cross-entropy (BCE) loss implementation of InfoMax objective (Hjelm et al., 2019; Yeh and Chen, 2019)³. In the current OOD detection problem, the loss is formulated as the approximation of MI between prototype vector of ID support set and target queries $I(C^{id}, Q)$:

$$I(C^{id}, Q) \geq \mathbb{E}_{\mathbb{P}}[\log F(C^{id}, x^{id})] + \mathbb{E}_{\mathbb{Q}}[\log(1 - F(C^{id}, x^{ood}))], \quad (3)$$

³For a theoretical justification on how binary cross-entropy (BCE) loss approximates Mutual Information (MI) between two random variables, including the alternatives, we refer reader to the prior works on investigating InfoMax objective for deep representation learning (Belghazi et al., 2018; Hjelm et al., 2019; Tschannen et al., 2020; Kong et al., 2020).

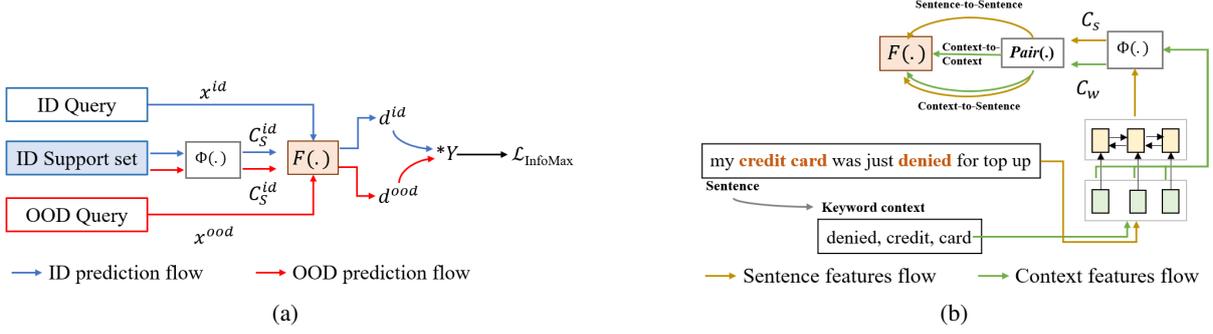


Figure 3: Proposed ProtoInfoMax. (a) ProtoInfoMax with prototype vector based on sentence features C_S . The encoder architecture that projects raw inputs into vectors in metric space is omitted to simplify the illustration. C_S^{id} is drew as separated ID and OOD lines to help the illustration, but in reality there is only one ID support set that is used to compute similarity scores d^{id} and d^{ood} . (b) A partial illustration of ProtoInfoMax++ with multiple prototype vectors (C_S, C_W), correspond to prototype vector based on sentence features and keyword context features respectively. Green boxes represent word embeddings in the encoder module of Prototypical Networks. Yellow boxes represent bidirectional GRU as sentence embedding layer.

where $\mathbb{E}_{\mathbb{P}}$ and $\mathbb{E}_{\mathbb{Q}}$ denote the expectation over ID and OOD samples respectively. x^{id} and x^{ood} are ID and OOD examples as target queries, $x^{id} \in Q^{id}$, $x^{ood} \in Q^{ood}$, $\{Q^{id}, Q^{ood}\} \in Q$. $F(\cdot)$ is a similarity scoring layer.

InfoMax loss is then defined as binary cross-entropy loss between ID and OOD prediction:

$$\mathcal{L}(C^{id}, Q) = \frac{1}{|Q^{id}|} \sum_{x^{id} \in Q^{id}} \log F(C^{id}, x^{id}) + \frac{1}{|Q^{ood}|} \sum_{x^{ood} \in Q^{ood}} \log(1 - F(C^{id}, x^{ood})). \quad (4)$$

4.2 Learning Framework

Figure 3a illustrates the proposed model with an InfoMax objective. The prediction outcome is represented by similarity scores between class prototypes and target queries, resulting in scores for ID targets (d^{id}) and OOD targets (d^{ood}). Since the training objective mainly focuses on promoting the separability between ID and OOD representations, we preserve ID supervision signals, i.e. $y \in \{\text{“positive”}, \text{“negative”}\}$ in sentiment classification benchmark, by projecting similarity scores (d^{id}, d^{ood}) into representation space of y . Thus, the final prediction is defined as: $\hat{y} = d * Y$; $d = \{d^{id}, d^{ood}\}$; $Y \in \mathbb{R}^{b \times d}$.

4.3 Prototype Generator $\Phi(\cdot)$

For both proposed models, we use standard prototype generator $\Phi(\cdot)$ based on feature averaging. Given encoded representations of ID support set $S^{id} \in \mathbb{R}^{b \times k \times d}$ (b = batch size, k = number of

examples in support set, d = dimension size of output representations), the prototype vector C^{id} is described as an averaged representation of those k -representations: $C^{id} = \Phi(S_k^{id}) = \frac{1}{k} \sum_{i=1}^k S_i^{id}$.

Sentence-based Features Given encoded representations of sentences in ID support set S^{id} , class prototype vector C^{id} is defined as a mean vector of those sentence features: $C^{id} = \Phi(S_k^{id})$. To prevent confusion, prototype vector based on sentence features is later denoted as C_S .

Keyword-based Features In an extremely low resource setting where training data may provide insufficient contexts due to the scarcity of novel sentences, the model may not be able to learn meaningful sentence representations. To better guide the learning, we utilize keywords as auxiliary inputs for ProtoInfoMax++ (Figure 3b).

Intuitively, sentences drawn from the same domain or intent distribution may share relevant context via their keywords. Therefore, keywords can be viewed as local context representation of a sentence. The more keywords that two sentences share in common, the more similar or related the two sentences are. While, from the perspective of word orientation in embedding space, keywords that are close together with respect to their angular distance are expected to carry similar semantic meaning. Sentences containing those similar subset of keywords can be considered to carry similar or related semantics. This motivates us to incorporate keyword representations into the current prototypical learning problem.

Prototype vector C_w is defined as a mean vec-

tor representation of sentence’s keywords $W = \{w_1, w_2, \dots, w_n\}$ weighted by their corresponding Idf value:

$$C_w = \frac{1}{n} \sum_{i=1}^n (w_i * Idf_i). \quad (5)$$

Since sentence inputs are composed of ID support set S^{id} , ID target queries Q^{id} and OOD target queries Q^{ood} , prototype vector based on keyword features can be further denoted as: C_w^{sup} representing keywords in ID support set; and C_w^Q representing keywords in target queries. For ID support set containing k -sentences, C_w is averaged over n -keywords and k -sentence features: $C_w^{sup} = \frac{1}{k} \frac{1}{n} \sum_{j=1}^k \sum_{i=1}^n (w_i * Idf_i)^j$.

4.4 Similarity function $F(\cdot)$

For model utilizing keyword auxiliary inputs (ProtoInfoMax++), we use a multi-perspective of similarity function $F(\cdot)$ to calculate similarity score d between support set and target queries.

Sentence-to-sentence similarity $F(C_S, Q)$ – This function is by default similarity measure for all Prototypical Networks in this study. Here, C_S denotes prototype vector of ID support set and Q is sentence embedding projection of target queries.

Context-to-context similarity $F(C_w^{sup}, C_w^Q)$ – We want to maximize MI between prototype representation of keywords in support set and target queries. C_w^{sup} is prototype vector computed from keyword contexts in support set, while C_w^Q is computed from keywords in target queries.

Context-to-sentence similarity $F(C_w^{sup*}, C_w^{Q*})$ – We want to maximize MI between sentences that share relevant context or keyword representations. Sentence representations with respect to their keyword contexts are computed as an element-wise matrix multiplication between encoded sentences and encoded keywords: $C_w^{sup*} = C_w^{sup} * C_S$; $C_w^{Q*} = C_w^Q * Q$.

4.5 Total Loss

ProtoInfoMax Given prototype vector based on sentence features C_S and target queries Q drawn from \mathcal{P}_{id} and \mathcal{P}_{ood} , the loss function for ProtoInfoMax is described as error loss given prototype vector generated from sentence features C_S and target queries Q :

$$\mathcal{L}_{\text{infomax}} = \mathcal{L}(C_S, Q). \quad (6)$$

ProtoInfoMax++ The total loss for ProtoInfoMax++ is described as cumulative losses given sentence-to-sentence similarity $F(C_S, Q)$, context-to-context similarity $F(C_w^{sup}, C_w^Q)$, and context-to-sentence similarity $F(C_w^{sup*}, C_w^{Q*})$ (section 4.4):

$$\begin{aligned} \mathcal{L}_{\text{infomax++}} = & \mathcal{L}(C_S, Q) + \\ & \mathcal{L}(C_w^{sup}, C_w^Q) + \\ & \mathcal{L}(C_w^{sup*}, C_w^{Q*}). \end{aligned} \quad (7)$$

5 Experiments

5.1 Dataset

Amazon Product Reviews For structuring Amazon review data into meta-tasks, we followed strategy from previous works on few-shot classification (Yu et al., 2018; Tan et al., 2019).

AI Conversational Data For constructing intent classification meta-tasks, we use two data sets that share contexts: AI Conversational Data (Chatterjee and Sengupta, 2020); and (CLINC150) (Larson et al., 2019; Casanueva et al., 2020)⁴. The preprocessed data contains disjoint classes across tasks, introducing a more challenging ID and OOD prediction task for Prototypical Networks in this study. In meta-training, each task (domain) is composed of 10 intent category labels ($N = 10$). Meta-validation and meta-testing are constructed from CLINC150. We use $N = 1$ and $N = 2$ set up to inspect model performance on one ID class and multiple ID classes prediction respectively.

5.2 Model and Hyper-parameters

Baselines We use two baselines: 1) **Proto-Net** (Snell et al., 2017; Yu et al., 2018), a native Prototypical Network with entropy-based loss function; 2) **O-Proto** (Tan et al., 2019), state-of-the-art approach for simultaneously learning ID classification and OOD detection. We do not include previous approaches based on non-Prototypical Networks (OSVM (Schölkopf et al., 2001), LSTM Autoencoder (Ryu et al., 2017), and vanilla CNN (Tan et al., 2019)) because these methods were shown to be under-performed in (Tan et al., 2019). We want to focus on further inspecting the reliability aspect of simple Prototypical Networks without additional learning pipelines.

⁴We use different benchmarks for intent classification task because the footage of preprocessed data from previous work (Tan et al., 2019) is unavailable publicly.

Data	Meta-training			Meta-validation			Meta-test		
	#Task	#ID class	#Sample	#Task	#ID sample	#OOD sample	#Task	#ID sample	#OOD sample
Amazon-rev (Sentiment)	13	2 (shared)	2M/task	4	(200,50)	20	4	(200,50)	20
AI-conv (Intent)	4	10 (disjoint)	100-3K/ID class	10	(120,30)	20	10	(120,30)	20

Table 1: Data set statistics. #ID sample in meta-validation and meta-test (n_A, n_B) shows the number of disjoint samples for ID support set and total samples for ID target query respectively. Except for meta-training in AI-conv data (Intent), number of samples (#Sample, #ID sample, #OOD sample) are shown as a figure representing examples within each task or domain.

Hyper-parameters For all models, we initialized word representation from pretrained fastText⁵. We updated fastText representation by further training it on current benchmark data, before using it as initialization for word embedding layer of Prototypical Networks. We used Tf-Idf (Sparck Jones, 1972; Salton and Buckley, 1988) as a keyword extraction method in the preparation of auxiliary inputs (keywords) for ProtoInfoMax++ due to its simple assumption. Namely, TfIdf measures word importance based on co-occurrence of words within a small group of documents. For future reference, this TfIdf approach can be substituted by any automated keyword extraction methods.

We use one layer Bidirectional-GRU (dimension size=200) as backbone encoder architecture for all models; and one layer Attention Network that is initialized based on r context query representations ($r = 5$) sampled from uniform distribution $\mathcal{U}[,1, .1]$. Similarity scoring layer $F(\cdot)$ is based on cosine similarity via matrix multiplication between prototype vector and target queries. All models were trained up to 60 epochs with batch size 100. Note that each epoch contains #Tasks that are dynamically sampled as training episodes.

5.3 Evaluation Metrics

ID and OOD Detection Errors We use (i) **Equal Error Rate (EER)** for measuring error in predicting OOD; (ii) **Class Error Rate (CER^{id})** for measuring error in predicting ID examples; and (iii) **CER^{all}** for measuring error in ID prediction given both ID and OOD subsets, following the previous work on OOD detection (Ryu et al., 2018; Tan et al., 2019). Except for CER^{id}, metrics are calculated based on heuristically selected threshold value τ . Given prediction outcomes with respect to decision whether examples are ID or OOD based on threshold τ , the error rate scores are defined as:

$$\text{FAR} = \frac{\text{FN}}{\# \text{ OOD examples}}, \quad (8)$$

⁵<https://fasttext.cc/>

$$\text{FRR} = \frac{\text{FP}}{\# \text{ ID examples}}, \quad (9)$$

$$\text{EER} = \frac{1 - (\text{TP} + \text{TN})}{\# \text{ Examples}}, \quad (10)$$

$$\text{CER}^{\text{id}} = \frac{\text{TP}^{\text{id}}}{\# \text{ ID examples}}, \quad (11)$$

$$\text{CER}^{\text{all}} = \frac{\text{TP}}{\# \text{ ID examples}}, \quad (12)$$

where TN denotes correct OOD prediction based on threshold τ . TP denotes correct ID prediction. FN measures OOD samples that are predicted as ID. FP measures ID samples that are predicted as OOD. TP^{id} is the number of correctly classified ID examples, excluding OOD samples.

Threshold score τ is calculated by heuristically searching a score conditioned by FRR and FAR metrics over sorted meta-test predictions (Ryu et al., 2018; Tan et al., 2019). That is, a score where the difference between False Acceptance Rate (FAR) and False Rejection Rate (FRR) has reached a minimum lower bound (FRR-FAR->0). Prior to the search, the initial threshold was defined as an average score of two prediction outcomes with the lowest scores. The final selected threshold τ is then used as a boundary score to distinguish between ID and OOD prediction.

Reliability Diagram Reliability diagram (Niculescu-Mizil and Caruana, 2005; Guo et al., 2017) depicts gaps between accuracy and model confidence. The larger the gap, the less calibrated the model is. That is, either the model is being underconfident or overconfident on estimating the winning predicted class labels. We use Expected

⁶Notice that our results (O-Proto performance) is different from those reported in (Tan et al., 2019). This might be due to different implementation frameworks: PyTorch vs. native Tensorflow; different hyper-parameters: we use $60 \times \# \text{ Task} \times 100$ batches due to our computational constraints vs. $5K \times \# \text{ Task} \times 100$ in (Tan et al., 2019); or different computing resources: GPU/CPU capacity used to train the models.

Method	Sentiment Cls ($N = 2$)						Intent Cls ($N = 1$)		Intent Cls ($N = 2$)		
	EER		CER ^{id}		CER ^{all}		EER		EER	CER ^{id}	CER ^{all}
	\mathcal{T}_{val}	\mathcal{T}_{test}	\mathcal{T}_{val}	\mathcal{T}_{test}	\mathcal{T}_{val}	\mathcal{T}_{test}	\mathcal{T}_{val}	\mathcal{T}_{test}			
Baselines											
Proto-Net (\mathcal{L}_{id})	0.398	0.387	0.266	0.285	0.445	0.536	0.456	0.420	0.525	0.316	0.634
O-Proto ($\mathcal{L}_{id}^{ent} + \mathcal{L}_{id}^{hinge} + \mathcal{L}_{ood}^{hinge}$)	0.348	0.375	0.411	0.409	0.631	0.643	0.404	0.390	0.482	0.373	0.683
This study											
ProtoInfoMax	0.373	0.278	0.351	0.365	0.592	0.521	0.398	0.368	0.398	0.256	0.549
ProtoInfoMax++	0.335	0.245	0.301	0.315	0.532	0.469	0.369	0.382	0.388	0.225	0.519

Table 2: Performance for $K = 100$ ⁶. The lower the better. Scores are based on top-3 the highest accuracy score for ID prediction (1-CER^{id}) across meta-validation and meta-test episodes (epochs). For one class prediction of intent classification ($N = 1$), EER and (1-CER^{all}) are equal, and CER^{id} = 1.0 because the number of ID class within the subset is 1. Evaluation for both $N = 1$ and $N = 2$ intent classification use the same model trained on $N = 10$, $K = 100$.

Calibration Error (ECE) (Naeini et al., 2015; Guo et al., 2017) to summarize the difference in expectation between confidence and accuracy (gaps) across all bins. We use similarity score d as a model confidence measure, following relevant work on distance-based prototypical learning (Xing et al., 2020).

6 Results and Analysis

We demonstrate the effectiveness of our proposed methods (**ProtoInfoMax** and **ProtoInfoMax++**) on two benchmarks for OOD detection (Table 2). Notice that native Prototypical Networks (**ProtoNet**) performs reasonably well, specifically for ID prediction (see scores based on CER^{id} and CER^{all}). However, this result can occur to models that always output predictions with a high score (e.g. high similarity score based on d in the current work), regardless whether the prediction is correct. The insight into this overconfidence behaviour is provided in Section 6.3 and 6.4.

6.1 Performance in different K-shot

Our **ProtoInfoMax** and **ProtoInfoMax++** also show a considerably consistent performance on meta-testing tasks under different K -shot values (Table 3 and 4), outperforming O-Proto.

6.2 On Threshold Score, FAR, and FRR

We want to further inspect the reliability of model prediction. Figure 4 shows the selected threshold score across models in the intent classification task. It can be observed that **O-Proto** has a tendency to be overconfident, suggested by a considerably high threshold score ($\tau = 0.97$ at epoch 0 and $\tau = 0.93$ at epoch 40). Both **ProtoInfoMax** and **ProtoInfoMax++** are being less confident after several epochs, yielding lower thresholds

Model	EER	CER ^{id}	CER ^{all}
K=1			
O-Proto	0.381	0.450	0.676
ProtoInfoMax	0.313	0.432	0.616
ProtoInfoMax++	0.335	0.430	0.615
K=10			
O-Proto	0.311	0.425	0.606
ProtoInfoMax	0.286	0.419	0.578
ProtoInfoMax++	0.254	0.375	0.537
K=100			
O-Proto	0.375	0.409	0.643
ProtoInfoMax	0.278	0.365	0.521
ProtoInfoMax++	0.245	0.315	0.469

Table 3: Performance under different K -shot values in sentiment classification ($N = 2$). Scores are based on the highest accuracy ($1 - \text{CER}^{\text{id}}$) on \mathcal{T}_{test} .

Model	EER	CER ^{id}	CER ^{all}
K=1			
O-Proto	0.515	0.391	0.698
ProtoInfoMax	0.480	0.397	0.674
ProtoInfoMax++	0.452	0.384	0.638
K=10			
O-Proto	0.493	0.402	0.694
ProtoInfoMax	0.451	0.400	0.686
ProtoInfoMax++	0.401	0.329	0.598
K=100			
O-Proto	0.482	0.373	0.683
ProtoInfoMax	0.398	0.256	0.549
ProtoInfoMax++	0.388	0.225	0.519

Table 4: Performance under different K -shot values in intent classification ($N = 2$).

($\tau = 0.87$ and $\tau = 0.74$ respectively). Compared to O-Proto, **ProtoInfoMax++** converges faster in early episodes (epoch= 0), yielding lower threshold score ($\tau = 79$).

We argue that one potential reason causing ProtoInfoMax++ to yield the lowest threshold score at epoch 0 is due to an effective regularization via InfoMax objective, in addition to the use of mul-

multiple representations that promotes discriminative metric space. Early convergence pattern can also be observed on ProtoInfoMax, indicating that InfoMax objective is empirically shown to be beneficial on preventing overconfidence from early iteration. Notice that the gaps between FAR and FRR for both **ProtoInfoMax** and **ProtoInfoMax++** at epoch = 40 are smaller. This indicates that both models underestimate ID and OOD samples, assigning them with low similarity scores ($d \leq 0.0$) with respect to ID class prototypes ⁷.

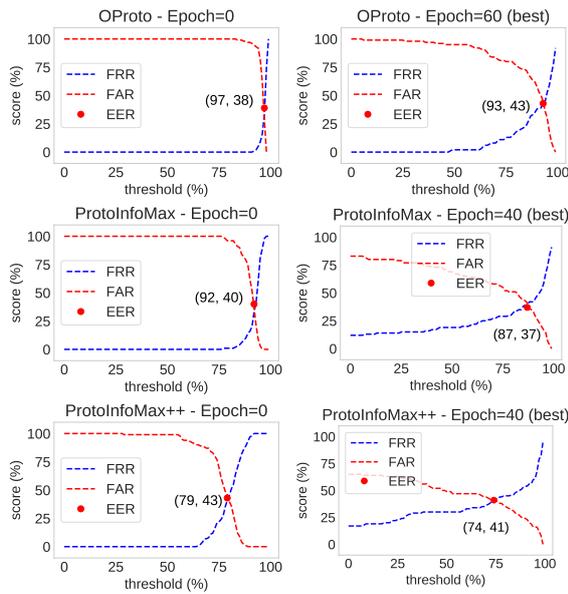


Figure 4: EER, FRR, FAR in intent classification meta-testing. Score (%) denotes proportion of samples that are either rejected (ID) or accepted (OOD) based on the selected threshold. To plot the above FAR and FRR, 200 prediction points corresponding to ID and OOD test samples were drawn randomly from 6 domains in $N = 2$ meta-test episodes.

6.3 Reliability in ID Prediction

Figure 5 compares the reliability of models in sentiment classification ⁸. In general, all models in this study tend to be overconfident, suggesting that future work focusing on directly tackling and investigating such problem is essential ⁹.

Compared to the baselines, our proposed **Pro-**

⁷We do not normalize $d, d \in [-1, 1]$ here to inspect whether the model penalizes OOD samples severely with similarity score $d \leq 0.0$.

⁸Since OOD labels are unknown during training, this evaluation only includes the prediction outcomes from ID target queries as test samples.

⁹In current work, we abuse terminology of “confidence score” to refer to similarity score d , following relevant work on distance-based prototypical learning (Xing et al., 2020).

toInfoMax and **ProtoInfoMax++** are shown to be less prone to typical overconfidence problem with respect to smaller gaps between their confidence score and the prediction accuracy. **Proto-Net**, however, suffers greatly from such overconfidence problem. It can be observed that Proto-Net assigns high similarity scores ($d \geq 0.9$) for all prediction points (see accuracy is lower than confidence score in Figure 5a).

Our methods achieve the lowest ECE scores (ECE ProtoInfoMax = 18.66 and ECE ProtoInfoMax++ = 16.40), suggesting a better reliability with respect to smaller gaps between model’s confidence score and prediction accuracy. **O-Proto** (Figure 5c) and **ProtoInfoMax** (Figure 5b) have both low confidence and overconfidence prediction. The models underestimate correct ID target queries (large gaps with high accuracy for $d \in (0.0, 0.2)$) and overestimate incorrect ID examples (large gaps with lower accuracy for $d \in (0.7, 1.0)$).

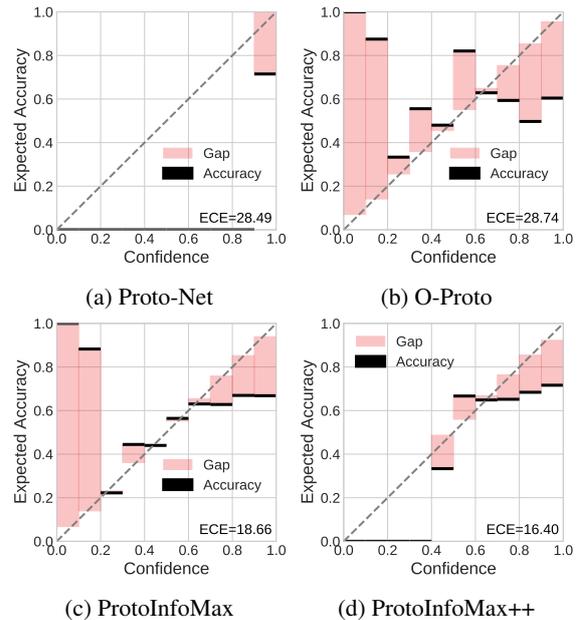


Figure 5: Reliability Diagram for ID prediction. Confidence scores were taken from \mathcal{T}^{test} in sentiment classification ($N = 2, K = 100$) based on the highest $1 - CER^{id}$.

6.4 Reliability in OOD Prediction

The reliability based on confidence histogram for ID and OOD prediction is provided in Figure 6 ¹⁰.

¹⁰Since OOD labels are unavailable during training, the reliability diagram is not applicable for evaluating OOD prediction.

In general, all models overestimate their prediction given ID target queries (see that the average confidence is higher than accuracy in Figure 6a, 6c, 6e). However, compared to **O-Proto**, our **ProtoInfoMax** and **ProtoInfoMax++** have a higher accuracy in ID classification tasks given their reasonably high confidence. Notice that for the ID prediction task, **ProtoInfoMax++** is more confident than the other two models ($d \in (0.4, 1.0)$) in Figure 6e).

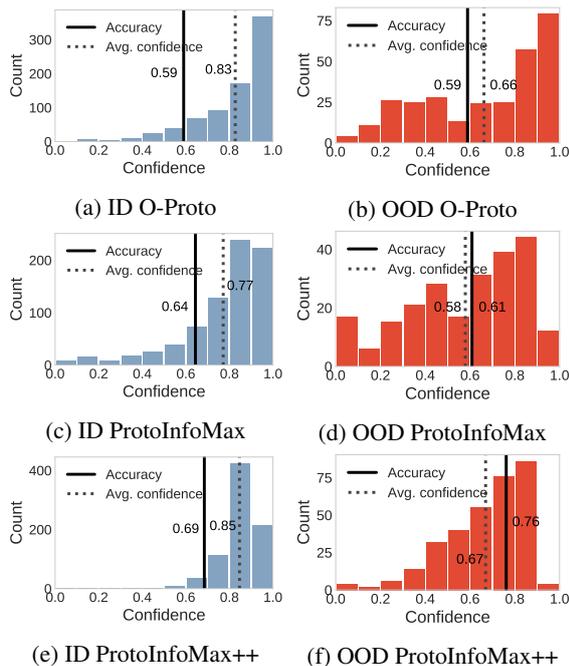


Figure 6: Confidence histogram for ID and OOD prediction. Confidence scores were taken from \mathcal{T}^{test} in sentiment classification ($N = 2, K = 100$). Note that the value of accuracy and average confidence here are not as precise as 1-CER^{id} score, since they were averaged across normalized bin scores.

For OOD detection ¹¹, our **ProtoInfoMax** and **ProtoInfoMax++** are shown to be less prone to overconfidence problem than **O-Proto**. See that the average confidence scores of both models are lower than their prediction accuracy (ProtoInfoMax avg. confidence: 0.58 in Figure 6d and ProtoInfoMax++: 0.67 in Figure 6f). In contrast, the average confidence score of **O-Proto** is higher than its prediction accuracy (Avg. Confidence = 0.66, Accuracy = 0.59 in Figure 6b), indicating the model prediction with an overconfidence issue.

¹¹Here, we view the task as one class OOD prediction where test samples contain OOD target queries only. Values below threshold τ are classified as OOD. Values above threshold are classified as ID.

7 Conclusion

Simultaneously learning In-Domain (ID) text classification and Out-of-Domain (OOD) detection under low resource constraints is realistic but under-explored. In this study, we aim at effectively and reliably learning zero-shot Out-of-Domain detection via Mutual Information Maximization (InfoMax) objective. Although we do not specifically tackle overconfidence problem of Neural Networks by calibrating models during training and evaluation stage in the current OOD detection task, we observe that the proposed ProtoInfoMax and ProtoInfoMax++ are less prone to such typical overconfidence problem compared to existing approaches. Overall, we improve performance of existing approaches up to 20% for OOD detection in low resource text classification.

Acknowledgment

This research is supported by Indonesian Endowment Fund for Education (LPDP) Scholarship under Beasiswa Pendidikan Indonesia (BPI) – ID Number 0003194/SC/D/9/LPDP2016. The content of the information does not necessarily reflect the position or the policy of the Government, and no official endorsement should be inferred.

References

- Yujia Bao, Menghua Wu, Shiyu Chang, and Regina Barzilay. 2020. Few-shot text classification with distributional signatures. In *International Conference on Learning Representations*.
- Mohamed Ishmael Belghazi, Aristide Baratin, Sai Rajeshwar, Sherjil Ozair, Yoshua Bengio, Aaron Courville, and Devon Hjelm. 2018. Mutual information neural estimation. In *International Conference on Machine Learning*, pages 531–540.
- Anthony J. Bell and Terrence J. Sejnowski. 1995. An Information-Maximization Approach to Blind Separation and Blind Deconvolution. *Neural Computation*, 7(6):1129–1159.
- Iñigo Casanueva, Tadas Temčinas, Daniela Gerz, Matthew Henderson, and Ivan Vulić. 2020. Efficient intent detection with dual sentence encoders. In *Proceedings of the 2nd Workshop on Natural Language Processing for Conversational AI*, pages 38–45, Online. Association for Computational Linguistics.
- Ajay Chatterjee and Shubhashis Sengupta. 2020. Intent mining from past conversations for conversational agent. In *Proceedings of the 28th International Conference on Computational Linguistics*,

- pages 4140–4152, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Thomas Dopierre, Christophe Gravier, and Wilfried Logerai. 2021. [A neural few-shot text classification reality check](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 935–943, Online. Association for Computational Linguistics.
- Meng Fang, Yuan Li, and Trevor Cohn. 2017. [Learning how to active learn: A deep reinforcement learning approach](#). In *Proceedings of the Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 595–605.
- Geli Fei and Bing Liu. 2016. [Breaking the closed world assumption in text classification](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 506–514, San Diego, California. Association for Computational Linguistics.
- Geli Fei, Shuai Wang, and Bing Liu. 2016. [Learning cumulatively to become more knowledgeable](#). In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining, KDD '16*, page 1565–1574, New York, NY, USA. Association for Computing Machinery.
- Tianyu Gao, Xu Han, Hao Zhu, Zhiyuan Liu, Peng Li, Maosong Sun, and Jie Zhou. 2019. [FewRel 2.0: Towards more challenging few-shot relation classification](#). 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)*, pages 6250–6255, Hong Kong, China. Association for Computational Linguistics.
- Chuanxing Geng, Sheng-jun Huang, and Songcan Chen. 2020. Recent advances in open set recognition: A survey. *IEEE Transactions on Pattern Analysis and Machine Intelligence*.
- Chuan Guo, Geoff Pleiss, Yu Sun, and Kilian Q. Weinberger. 2017. On calibration of modern neural networks. In *Proceedings of the 34th International Conference on Machine Learning - Volume 70, ICML'17*, page 1321–1330. JMLR.org.
- Xu Han, Hao Zhu, Pengfei Yu, Ziyun Wang, Yuan Yao, Zhiyuan Liu, and Maosong Sun. 2018. [FewRel: A large-scale supervised few-shot relation classification dataset with state-of-the-art evaluation](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4803–4809, Brussels, Belgium. Association for Computational Linguistics.
- Dan Hendrycks, Mantas Mazeika, and Thomas Dietterich. 2019. Deep anomaly detection with outlier exposure. *Proceedings of the International Conference on Learning Representations*.
- Devon Hjelm, Alex Fedorov, Samuel Lavoie-Marchildon, Karan Grewal, Philip Bachman, Adam Trischler, and Yoshua Bengio. 2019. [Learning deep representations by mutual information estimation and maximization](#). In *ICLR 2019*. ICLR.
- Shehroz S Khan and Michael G Madden. 2014. One-class classification: taxonomy of study and review of techniques. *The Knowledge Engineering Review*, 29(3):345–374.
- Joo-Kyung Kim and Young-Bum Kim. 2018a. Joint learning of domain classification and out-of-domain detection with dynamic class weighting for satisfying false acceptance rates. *Proc. Interspeech 2018*, pages 556–560.
- Joo-Kyung Kim and Young-Bum Kim. 2018b. [Joint learning of domain classification and out-of-domain detection with dynamic class weighting for satisfying false acceptance rates](#). In *Proc. Interspeech 2018*, pages 556–560.
- Lingpeng Kong, Cyprien de Masson d’Autume, Lei Yu, Wang Ling, Zihang Dai, and Dani Yogatama. 2020. [A mutual information maximization perspective of language representation learning](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.
- Balaji Lakshminarayanan, Alexander Pritzel, and Charles Blundell. 2017. Simple and scalable predictive uncertainty estimation using deep ensembles. In *Proceedings of the 31st International Conference on Neural Information Processing Systems, NIPS'17*, page 6405–6416, Red Hook, NY, USA. Curran Associates Inc.
- Ian Lane, Tatsuya Kawahara, Tomoko Matsui, and Satoshi Nakamura. 2006. Out-of-domain utterance detection using classification confidences of multiple topics. *IEEE Transactions on Audio, Speech, and Language Processing*, 15(1):150–161.
- Stefan Larson, Anish Mahendran, Joseph J. Peper, Christopher Clarke, Andrew Lee, Parker Hill, Jonathan K. Kummerfeld, Kevin Leach, Michael A. Laurenzano, Lingjia Tang, and Jason Mars. 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)*, pages 1311–1316, Hong Kong, China. Association for Computational Linguistics.
- Mike Lewis and Angela Fan. 2019. [Generative question answering: Learning to answer the whole question](#). In *International Conference on Learning Representations*.
- Jiwei Li, Michel Galley, Chris Brockett, Jianfeng Gao, and Bill Dolan. 2016. [A diversity-promoting objective function for neural conversation models](#). In

- Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 110–119, San Diego, California. Association for Computational Linguistics.
- Shiyu Liang, Yixuan Li, and R Srikant. 2018. Enhancing the reliability of out-of-distribution image detection in neural networks. In *International Conference on Learning Representations*.
- R. Linsker. 1988. [Self-organization in a perceptual network](#). *Computer*, 21(3):105–117.
- Mahdi Pakdaman Naeni, Gregory Cooper, and Milos Hauskrecht. 2015. Obtaining well calibrated probabilities using bayesian binning. In *Twenty-Ninth AAAI Conference on Artificial Intelligence*.
- Alexandru Niculescu-Mizil and Rich Caruana. 2005. Predicting good probabilities with supervised learning. In *Proceedings of the 22nd international conference on Machine learning*, pages 625–632.
- Pramuditha Perera, Ramesh Nallapati, and Bing Xiang. 2019. Ocgan: One-class novelty detection using gans with constrained latent representations. In *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Mirco Ravanelli and Yoshua Bengio. 2019. Learning speaker representations with mutual information. *Proc. Interspeech 2019*, pages 1153–1157.
- Lukas Ruff, Robert Vandermeulen, Nico Goernitz, Lucas Deecke, Shoaib Ahmed Siddiqui, Alexander Binder, Emmanuel Müller, and Marius Kloft. 2018. [Deep one-class classification](#). volume 80 of *Proceedings of Machine Learning Research*, pages 4393–4402, Stockholmsmässan, Stockholm Sweden. PMLR.
- Seonghan Ryu, Seokhwan Kim, Junhwi Choi, Hwanjo Yu, and Gary Geunbae Lee. 2017. [Neural sentence embedding using only in-domain sentences for out-of-domain sentence detection in dialog systems](#). *Pattern Recognition Letters*, 88:26 – 32.
- Seonghan Ryu, Sangjun Koo, Hwanjo Yu, and Gary Geunbae Lee. 2018. [Out-of-domain detection based on generative adversarial network](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 714–718, Brussels, Belgium. Association for Computational Linguistics.
- Gerard Salton and Christopher Buckley. 1988. Term-weighting approaches in automatic text retrieval. *Information processing & management*, 24(5):513–523.
- W. J. Scheirer, A. de Rezende Rocha, A. Sapkota, and T. E. Boult. 2013. [Toward open set recognition](#). *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(7):1757–1772.
- Bernhard Schölkopf, John C. Platt, John Shawe-Taylor, Alex J. Smola, and Robert C. Williamson. 2001. [Estimating the Support of a High-Dimensional Distribution](#). *Neural Computation*, 13(7):1443–1471.
- Alireza Shafaei, Mark Schmidt, and James Little. 2019. A Less Biased Evaluation of Out-of-distribution Sample Detectors. In *British Machine Vision Conference (BMVC)*.
- Lei Shu, Hu Xu, and Bing Liu. 2017. [DOC: Deep open classification of text documents](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 2911–2916, Copenhagen, Denmark. Association for Computational Linguistics.
- Jake Snell, Kevin Swersky, and Richard Zemel. 2017. Prototypical networks for few-shot learning. In *Advances in neural information processing systems*, pages 4077–4087.
- Karen Sparck Jones. 1972. A statistical interpretation of term specificity and its application in retrieval. *Journal of Documentation*, 28(1):11–21.
- Shengli Sun, Qingfeng Sun, Kevin Zhou, and Tengchao Lv. 2019. [Hierarchical attention prototypical networks for few-shot text classification](#). 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)*, pages 476–485, Hong Kong, China. Association for Computational Linguistics.
- Ming Tan, Yang Yu, Haoyu Wang, Dakuo Wang, Saroni Potdar, Shiyu Chang, and Mo Yu. 2019. [Out-of-domain detection for low-resource text classification tasks](#). 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)*, pages 3566–3572, Hong Kong, China. Association for Computational Linguistics.
- Michael Tobias Tschannen, Josip Djolonga, Paul Kishan Rubenstein, Sylvain Gelly, and Mario Lučić. 2020. [On mutual information maximization for representation learning](#). In *International Conference on Learning Representations*. Michael Tschannen and Josip Djolonga contributed equally.
- Gokhan Tur, Anoop Deoras, and Dilek Hakkani-Tür. 2014. Detecting out-of-domain utterances addressed to a virtual personal assistant. In *Fifteenth Annual Conference of the International Speech Communication Association*.
- Chen Xing, Sercan Arik, Zizhao Zhang, and Tomas Pfister. 2020. [Distance-based learning from errors for confidence calibration](#). In *International Conference on Learning Representations*.

- Yi-Ting Yeh and Yun-Nung Chen. 2019. [QAInfomax: Learning robust question answering system by mutual information maximization](#). 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)*, pages 3370–3375, Hong Kong, China. Association for Computational Linguistics.
- Mo Yu, Xiaoxiao Guo, Jinfeng Yi, Shiyu Chang, Saloni Potdar, Yu Cheng, Gerald Tesauro, Haoyu Wang, and Bowen Zhou. 2018. [Diverse few-shot text classification with multiple metrics](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 1206–1215, New Orleans, Louisiana. Association for Computational Linguistics.
- Houssam Zenati, Manon Romain, Chuan-Sheng Foo, Bruno Lecouat, and Vijay Chandrasekhar. 2018. Adversarially learned anomaly detection. In *2018 IEEE International Conference on Data Mining (ICDM)*, pages 727–736. IEEE.
- Jianguo Zhang, Kazuma Hashimoto, Wenhao Liu, Chien-Sheng Wu, Yao Wan, Philip Yu, Richard Socher, and Caiming Xiong. 2020. [Discriminative nearest neighbor few-shot intent detection by transferring natural language inference](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5064–5082, Online. Association for Computational Linguistics.
- Yinhe Zheng, Guanyi Chen, and Minlie Huang. 2020. Out-of-domain detection for natural language understanding in dialog systems. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 28:1198–1209.