

IDAS: Intent Discovery with Abstractive Summarization

Maarten De Raedt^{✦♣} Frédéric Godin[✦] Thomas Demeester[♣] Chris Develder[♣]
[✦] Sinch Chatlayer [♣] Ghent University
{maarten.deraedt, thomas.demeester, chris.develder}@ugent.be
frederic.godin@sinch.com

Abstract

Intent discovery is the task of inferring latent intents from a set of unlabeled utterances, and is a useful step towards the efficient creation of new conversational agents. We show that recent competitive methods in intent discovery can be outperformed by clustering utterances based on *abstractive summaries*, i.e., “labels”, that retain the core elements while removing non-essential information. We contribute the IDAS approach, which collects a set of descriptive utterance labels by prompting a Large Language Model, starting from a well-chosen seed set of prototypical utterances, to bootstrap an In-Context Learning procedure to generate labels for non-prototypical utterances. The utterances and their resulting noisy labels are then encoded by a *frozen* pre-trained encoder, and subsequently clustered to recover the latent intents. For the *unsupervised* task (without any intent labels) IDAS outperforms the state-of-the-art by up to +7.42% in standard cluster metrics for the Banking, StackOverflow, and Transport datasets. For the *semi-supervised* task (with labels for a subset of intents) IDAS surpasses 2 recent methods on the CLINC benchmark without even using labeled data.

1 Introduction

Intent classification is ubiquitous in conversational modelling. To that end, finetuning Large Language Models (LLMs) on task-specific intent data has been proven very effective (Casanueva et al., 2020; Zhang et al., 2021d). However, such finetuning requires manually annotated (utterance, intent) pairs as training data, which are time-consuming and thus expensive to acquire. Companies often have an abundance of utterances relevant to the application area of their interest, e.g., those exchanged between customers and support agents, but manually annotating them remains costly. Consequently, intent discovery aims to recover latent intents without using any such manually annotated utterances, by partitioning a given set of (unlabeled) utterances into

Utterance	Generated label
find out when my next upcoming payday will be my next paycheck is available when what is the date of my last paycheck	when is next payday when is next payday when was last payday
i want to know how to change my oil what is the way to change motor oil how easy is it to change your own oil	how to change oil how to change oil DIY oil change
can you tell me the <i>apr</i> on my visa card what’s the annual rate on my discover card	<i>interest rate inquiry</i> interest rate inquiry

Table 1: *Illustration* based on GPT-3 and CLINC (Larson et al., 2019), demonstrating how *abstractly* summarizing utterances retains the core elements while removing non-intent related information. The example in the bottom block, where *apr* is labeled as *interest rate inquiry*, exemplifies the broad domain knowledge captured by LLMs.

clusters, where utterances within a cluster should share the same *conversational goal* or *intent*.

Prior works typically (i) train an unsupervised sentence encoder to map utterances to vectors, after which these are (ii) clustered to infer latent intents. Such unsupervised encoder training is achieved largely under the assumption that utterances with similar encodings convey the same intent. For instance, by iteratively clustering and updating the encoder with supervision from the cluster assignments (Xie et al., 2016a; Caron et al., 2018a; Hadifar et al., 2019; Zhang et al., 2021c), or by retrieving utterances with similar encodings and using them as positive pairs to train the encoder with contrastive learning (Zhang et al., 2021a, 2022).

Yet, it remains unclear which particular features cause utterance representations to be similar. Various noisy features unrelated to the underlying intents, e.g., *syntax*, *n-gram overlap*, *nouns*, etc. may contribute in making utterances similar, leading to sentence encoders whose vector encodings may inadequately represent the underlying intents.

Different from prior works that train unsupervised encoders, we use a pre-trained encoder without requiring any further finetuning, since we pro-

pose making utterances more (dis)similar in the textual space by *abstractly* summarizing them into concise descriptions, i.e., “labels”, that preserve their core elements while removing non-essential information. We hypothesize that these core elements better represent intents and prevent non-intent related information from influencing the vector similarity. Table 1 illustrates how labels retain the intent-related information by discarding irrelevant aspects such as syntax and nouns.

This paper introduces Intent Discovery with Abstractive Summarization (IDAS in short), whereby the label generation process builds upon recent advancements of In-Context Learning (ICL) (Brown et al., 2020). In ICL, an LLM is prompted with an instruction including a small number of (input, output) demonstrations of the task at hand. ICL has shown to be effective at few-shot learning without additional LLM finetuning (Min et al., 2022a,b). However, intent discovery is unsupervised and therefore lacks the annotated (utterance, label) demonstrations required for ICL. To overcome this limitation, our proposed IDAS proceeds in four steps. First, a subset of diverse *prototypical* utterances representative of distinct latent intents are identified by performing an initial clustering and selecting those utterances closest to each cluster’s centroid, for which an LLM is then prompted to generate a short descriptive label. Second, labels for the remaining *non-prototypical* utterances are obtained by retrieving the subset of the n utterances most similar to the input utterance, from the continually expanding set of utterances with already generated labels (initialized with just the prototypes), and using those n neighbors as ICL-demonstrations to generate the input utterance’s label. Third, as the generated labels may still turn out too general or noisy, utterances with their labels are combined into a single vector representation using a *frozen* pre-trained encoder. Finally, K -means clusters the combined encodings to infer latent intents.

We compare our IDAS approach with the state-of-the-art in unsupervised intent discovery on Banking (Casanueva et al., 2020), StackOverflow (Xu et al., 2015), and a private dataset from a transport company, to assess IDAS’s effectiveness in practice. We show that IDAS substantially outperforms the state-of-the-art, with average improvements in cluster metrics of +3.94%, +2.86%, and +3.34% in Adjusted Rand Index, Normalized Mutual Information, and Cluster Accuracy, respectively. Fur-

ther, IDAS surpasses two *semi-supervised* intent discovery methods on CLINC (Larson et al., 2019) despite not using any ground truth annotations.

2 Related Work

Statistical approaches: Early, more general short text clustering methods employ statistical methods such as tf-idf (Sparck Jones, 1972), to map text to vectors. Yet, the sparsity of these encodings prevents similar texts, but phrased with different synonyms, from being assigned to the same cluster. To specifically mitigate this synonym effect, external features have been used to enrich such *sparse* vectors, e.g., with WordNet (Miller, 1995) synonyms or lexical chains (Hotho et al., 2003; Wei et al., 2015), or Wikipedia titles or categories (Banerjee et al., 2007; Hu et al., 2009).

Neural sentence encoders: Rather than relying on external knowledge sources, neural approaches pre-train sentence encoders in a self-supervised way (Kiros et al., 2015; Gao et al., 2021), or with supervision (Conneau et al., 2017; Reimers and Gurevych, 2019; Gao et al., 2021), to produce *dense* general-purpose vectors that better capture synonymy and semantic relatedness.

Unsupervised intent discovery: Since general-purpose neural encoders may fail to capture domain-specific intent information, intent discovery solutions have shifted towards unsupervised sentence encoders specifically trained on the domain data at hand. For instance, Xu et al. (2015) train a self-supervised Convolutional Neural Network, and use it to encode and cluster utterances with K -means. Zhang et al. (2022) adopt the same 2-step approach, but instead pre-train the encoder with contrastive learning, where utterances with similar vector encodings are retrieved to serve as positive pairs. A more common strategy is to cluster and train the encoder end-to-end, either by (i) *iteratively* clustering utterances and updating the encoder with supervision from the cluster assignments (Xie et al., 2016a; Caron et al., 2018b; Hadifar et al., 2019), or (ii) *simultaneously* clustering utterances and updating the encoder’s weights with a joint loss criterion (Yang et al., 2017a; Zhang et al., 2021a).

As an alternative strategy to make utterances more (dis)similar based on the intents they convey, we employ an LLM to summarize utterances into labels that retain both the utterances’ core ele-

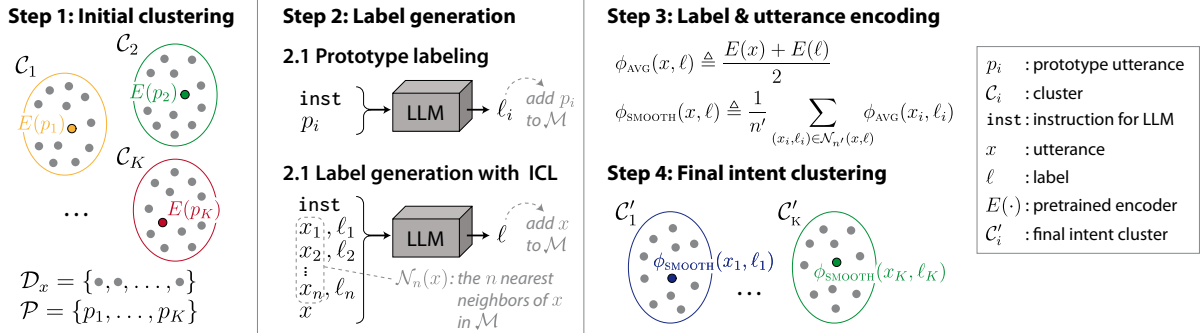


Fig. 1: Overview of our IDAS approach.

ments and domain-specific information as encoded in the LLM’s weights. Since our generated labels should increase the (dis)similarity of (un)related utterances in the input space, rather than directly in the vector space, we use a *frozen* pre-trained encoder, thus deviating from the above methods that *train* unsupervised encoders.

Semi-supervised intent discovery: Similar to our current work, the aforementioned methods focus on *unsupervised* intent discovery. In the related but different *semi-supervised* intent discovery task, a fraction of the latent intents is assumed to be known, i.e., the “Known Class Ratio”. Annotated data from these known intents is exploited to improve the detection of both known and unknown intent utterances, e.g., by optimizing a cluster loss with pairwise constraints derived from utterances of the same known intent (Lin et al., 2020). Alternative 2-step approaches first pre-train encoders with supervision from known intent utterances, then either directly encode and cluster utterances with K -means (Shen et al., 2021), or further refine the encoder on the unlabeled utterances. The latter refinement can be achieved through contrastive learning (Zhang et al., 2022) or by iteratively clustering and updating the encoder (Zhang et al., 2021b,c).

In-context learning: The core idea of ICL (Brown et al., 2020) is to perform tasks through inference, i.e., without updating parameters, by prompting an LLM with the string concatenation comprising (i) a task instruction, (ii) a small set of (input, output) demonstrations, and (iii) the input. We implement IDAS’s label generation process with ICL, as it has shown to substantially outperform zero-shot approaches *without* demonstrations (Min et al., 2022a,b; Chen et al., 2022). However, since we focus on unsupervised intent discovery and thus lack annotated (utterance, label) demon-

strations, we bootstrap the set of demonstrations with automatically retrieved “prototypes”. Rather than selecting demonstrations randomly, Liu et al. (2022) found that it is more effective to pick demonstrations similar to the input utterance, which we thus do. Note that alternative methods are possible (Rubin et al., 2022; Sorensen et al., 2022).

3 Methodology

Task formulation: Let $\{(x_i, y_i) | i = 1 \dots N\}$ be a dataset of N utterances $x \in \mathcal{X}$ from the set of natural language expressions \mathcal{X} , with corresponding intents y chosen from a set of K possible intents $\mathcal{Y} = \{y_i | i = 1 \dots K\}$. Given the utterances without the intents, $\mathcal{D}_x = \{x_i | i = 1 \dots N\}$, intent discovery aims to infer \mathcal{Y} from \mathcal{D}_x by mapping utterances x_i to vectors $E(x_i)$ with encoder $E : \mathcal{X} \rightarrow \mathbb{R}^d$, based on which the utterances are partitioned into clusters $\{\mathcal{C}_i | i = 1 \dots K\}$, such that clustered utterances (e.g., $x_{i,j}, x_{k,j} \in \mathcal{C}_j$) share the same intent ($y_{i,j} = y_{k,j}$), while utterances from different clusters (e.g., $x_{i,j} \in \mathcal{C}_j$ and $x_{k,l} \in \mathcal{C}_l$, $\mathcal{C}_l \neq \mathcal{C}_k$) have distinct intents ($y_{i,j} \neq y_{k,l}$).

Overview: As summarized in Fig. 1, to infer latent intents IDAS (1) identifies a subset of diverse “prototypes”, $\mathcal{P} \subset \mathcal{D}_x$, representative of the latent intents (§3.1); then (2) independently summarizes them into labels, which are further used to also generate labels for the remaining non-prototypical utterances $x \in \mathcal{D}_x \setminus \mathcal{P}$, by retrieving from the subset \mathcal{M} of utterances that already have labels (initially \mathcal{P}) the set $\mathcal{N}_n(x)$ of n utterances most similar to x as ICL-demonstrations for generating the label of x (§3.2); further (3) encodes utterances and their labels into a single vector representation with a *frozen* pre-trained encoder (§3.3); and finally (4) infers the latent intents by performing K -means on the combined representations (§3.4).

3.1 Step 1: Initial Clustering

The objective of this step is to identify a diverse set of prototypes, $\mathcal{P} \subset \mathcal{D}_x$, that in Step 2 will be automatically labeled by an LLM and serve as initial demonstrations for generating the labels of non-prototypical utterances. It is therefore important to choose prototypes $p \in \mathcal{P}$ that each represent a distinct latent intent $y \in \mathcal{Y}$, and collectively cover as many as possible of all latent intents. We assume a similarity function between two vector representations of utterances by $s : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$, and use it to retrieve prototypes by performing an initial clustering on the utterances in \mathcal{D}_x , in the vector representation space induced by encoder E . Then we select a prototype from each identified cluster, as the utterance in that cluster whose vector representation is closest to the cluster’s centroid.

Formally, the utterances in \mathcal{D}_x are first encoded with E and then partitioned into K ($=|\mathcal{Y}|$) clusters

$$\mathcal{C}_1, \dots, \mathcal{C}_K = K\text{-means}(\mathcal{D}_x),$$

for which the respective centroids $c_i \in \mathbb{R}^d$ and prototypes $p_i \in \mathcal{D}_x$ are calculated as

$$c_i = \frac{1}{|\mathcal{C}_i|} \sum_{x \in \mathcal{C}_i} E(x), \quad p_i = \operatorname{argmax}_{x \in \mathcal{C}_i} s(E(x), c_i).$$

3.2 Step 2: Label Generation

Step 2.1: Prototype Labeling To generate label ℓ_i for prototype p_i , we employ an LLM and provide it with an instruction (inst) such as “describe the question in a maximum of 5 words”. The LLM then generates a concise description of the prototype p_i , which we use as its label ℓ_i . Mathematically, this is represented as

$$\ell_i = \operatorname{argmax}_{\ell \in \mathcal{X}} P(\ell | \text{inst}, p_i),$$

where P denotes the probability distribution of the LLM, and ℓ_i represents the token sequence t_{1_i}, \dots, t_{i_i} output by the LLM.

Step 2.2: Label Generation with ICL To generate label ℓ for the non-prototypical utterance $x \in \mathcal{D}_x \setminus \mathcal{P}$, IDAS utilizes ICL by conditioning an LLM on the prompt, i.e., the string concatenation of (i) an instruction inst, e.g., “classify the question into one of the labels”, (ii) the set of n demonstrations of (utterance, label) pairs $\{(x_i, \ell_i) | i = 1 \dots n\}$, and (iii) the utterance x itself. Formally, the label is the token sequence generated by the LLM that maximizes the probability

given the prompt:

$$\ell = \operatorname{argmax}_{\ell \in \mathcal{X}} P(\ell | \text{inst}, x_1, \ell_1, \dots, x_n, \ell_n, x).$$

Since unsupervised intent discovery lacks manually annotated demonstrations, IDAS uses a continually expanding set of utterances with *automatically* generated labels, denoted by \mathcal{M} . Initially, $\mathcal{M} = \mathcal{P}$, with \mathcal{P} the set of prototypes from Step 2.1. An utterance x with newly generated label ℓ is added to \mathcal{M} , such that it can serve as a demonstration for remaining unlabeled utterances.

Typically, ICL uses a small set of n demonstrations (i) due to the limit on the number of input tokens of LLMs, and (ii) because performance does not improve for larger number of demonstrations (Min et al., 2022c). Moreover, Liu et al. (2022) found that selecting demonstrations as samples similar to the test input, rather than choosing them randomly, substantially boosts ICL’s performance. Therefore, IDAS adopts KATE (Liu et al., 2022) by first mapping utterances in \mathcal{M} to vectors with encoder E , and then using the similarity function s to select the set of the n most similar utterances¹ from \mathcal{M} to $E(x)$, denoted by $\mathcal{N}_n(x) \subset \mathcal{M}$, as demonstrations for input utterance x .

Note that while we use “classify” in the instruction, we do not consider the prototypical labels generated in Step 1 as a fixed label set (i.e., *verbalizers*). Rather, label ℓ for non-prototypical utterance x is the token sequence as generated directly by the LLM. As a result, labels for non-prototypical utterances may still differ from those generated for the prototypes. Particularly, the LLM can generate new labels for input utterances that represent intents for which no prototypes have been identified yet, and thus have no ICL demonstrations of the latent intent. Thus, we minimize error propagation from Step 1. On the other hand, when the LLM considers that a demonstration likely shares the same latent intent with the input utterance, the “classify” instruction should encourage the LLM to generate a copy of that demonstration’s label, which in turn minimizes variation among generated labels of utterances with the same latent intent.

3.3 Step 3: Encoding Utterances and Labels

After Step 2, each utterance $x \in \mathcal{D}_x$ has an associated generated label $\ell \in \mathcal{M}$. We use the pre-trained

¹We set hyperparameter n to 8, based on the findings of Min et al. (2022c); Lyu et al. (2022). Ablations for different n values are presented in §5.2.

encoder E to respectively encode the utterances and their corresponding labels into separate vectors $E(x)$ and $E(\ell)$, after which these are averaged into the combined representation:

$$\phi_{\text{AVG}}(x, \ell) \triangleq \frac{E(x) + E(\ell)}{2}. \quad (1)$$

(Note that utterances could also be represented just by their label encoding $E(\ell)$, yet such generated labels could be noisy or overly general.)

We further contribute a non-parametric smoothing method that (i) aims to suppress features that are specific to individual utterances and thus potentially less representative of the underlying intents, while (ii) enhancing those features that are shared across utterances and thus more likely to be representative of the latent intents. We therefore represent utterance x as the average of the vector encodings of the n' most similar utterances $\mathcal{N}_{n'}(x, \ell)$ to x , including x itself:

$$\phi_{\text{SMOOTH}}(x, \ell) \triangleq \frac{1}{n'} \sum_{(x_i, \ell_i) \in \mathcal{N}_{n'}(x, \ell)} \phi_{\text{AVG}}(x_i, \ell_i). \quad (2)$$

We automatically determine the value of n' as the value that maximizes the average silhouette score (Rousseeuw, 1987) among all samples, which for sample i is given by

$$\text{silhouette-score}(i) = \frac{b(i) - a(i)}{\max(a(i), b(i))},$$

where $a(i)$ is the average distance of sample i to all other samples in its cluster, and $b(i)$ is the average distance of sample i to all samples in the neighboring cluster nearest to i .

3.4 Step 4: Final intent discovery

To finally infer the latent intents, we represent each utterance $x \in \mathcal{D}_x$ with its label ℓ as $\phi_{\text{SMOOTH}}(x, \ell)$, and apply K -means clustering, setting K to the ground truth number of latent intents $|\mathcal{Y}|$, following Hadifar et al. (2019); Zhang et al. (2021a,c, 2022).

4 Experimental Setup

4.1 Datasets

We evaluate our IDAS approach on two widely adopted intent classification datasets, CLINC (Larson et al., 2019) and Banking (Casanueva et al., 2020), as well as the StackOverflow topic classification dataset (Xu et al., 2015). We also use a private dataset from a transportation company. Table 2 summarizes dataset statistics.

Dataset	# Train	# Test	# Intents
CLINC	18,000	2,250	150
Banking	9,016	3,080	77
Transport	-	1,257	42
StackOverflow	18,000	1,000	20

Table 2: Dataset statistics.

4.2 Baselines

On Banking, StackOverflow, and our Transport dataset, we compare IDAS against the state-of-the-art in unsupervised intent discovery, i.e., the MTP-CLNN model (Zhang et al., 2022) that outperforms prior unsupervised methods, such as DEC (Xie et al., 2016b), DCN (Yang et al., 2017b), and DeepCluster (Caron et al., 2018b). As the MTP-CLNN model is pre-trained on the annotated training data of CLINC, directly comparing against it would be unfair. Instead, we compare our approach on CLINC with two state-of-the-art *semi-supervised* intent discovery methods, DAC (Zhang et al., 2021c) and SCL+PLT (Shen et al., 2021). Compared to the semi-supervised setting, the unsupervised setting without annotations is thus more challenging. We report results of DAC and SCL+PLT with an increasing ‘‘Known Class Ratio’’ (KCR) of 25%, 50%, and 75%, using the annotated data for the known intents of Shen et al. (2021).

4.3 Evaluation

Following Zhang et al. (2021c); Shen et al. (2021); Zhang et al. (2022), we assess cluster performance by comparing the predicted clusters to the ground truth intents using the (i) Adjusted Rand Index (ARI) (Steinley, 2004), (ii) Normalized Mutual Information (NMI), and (iii) Cluster Accuracy (ACC) based on the Hungarian algorithm (Kuhn, 1955). Since IDAS’s label generation process may depend on the order in which utterances occur, we perform Steps 1–2 leading to utterance labels 5 times, shuffling the utterance order. We further conduct the final clustering Step 4 with 10 different seeds for each of those 5 label generation runs, to account for variation incurred by K -means. For each dataset, we then average the results in terms of means and standard variations across each of these 5 sets.

4.4 Implementation

Encoder: We use the same pre-trained encoder E in all steps of our approach, i.e., to (i) retrieve prototypes (§3.1), (ii) mine the n demonstrations

$\mathcal{N}_n(x)$ for utterance x (§3.2), and (iii) encode utterances with their labels using Eqs. (1)–(2) (§3.3). To rule out performance differences stemming purely from the encoder, we employ the same pre-trained encoder as the baseline we compare with: we use the MTP encoder for Banking, StackOverflow, and Transport, where we compare to MTP-CLNN (Zhang et al., 2022), and the SBERT encoder paraphrase-mpnet-base2 (i.e., SMPNET) (Reimers and Gurevych, 2019) for CLINC, where we compare to DAC (Zhang et al., 2021c) and SCL+PLT (Shen et al., 2021).

Language models and prompts: IDAS uses the text-davinci-003 GPT-3 model (Ouyang et al., 2022) as its LLM for label generation. We adopt the OpenAI playground default values, except for the temperature, which we set to 0 to minimize variation among generated labels of utterances with the same latent intent. To generate prototypical labels (§3.2), we use the instruction “Describe the domain question in a maximum of 5 words”, where the domain is *banking*, *chatbot*, or *transport* for the corresponding dataset. Since StackOverflow is a topic rather than an intent classification dataset, we adopt a slightly different prototypical prompt. To generate labels for non-prototypical utterances with ICL (§3.2), we use “Classify the domain question into one of the provided labels” for all 4 datasets. See Appendix A.2 for full prompts and examples.

Nearest neighbor retrieval: The function s is implemented with cosine similarity. We use $n=8$ demonstrations $\mathcal{N}_n(x)$ to generate label ℓ for utterance x (§3.2), based on Min et al. (2022c) and Lyu et al. (2022), who report that further increasing n does not improve ICL’s performance. The number of smoothing samples n' is determined by running the final K -means (§3.4) multiple times with n' ranging from 5 to 45 and selecting the value that maximizes the average silhouette score.

5 Results and Discussion

5.1 Main Results

In *unsupervised* clustering, no labels are available and thus there is only a test set, used to evaluate the model’s induced clusters against gold standard labels (Xie et al., 2016a; Yang et al., 2017a; Hadifar et al., 2019; Zhang et al., 2021a). In the *semi-supervised* intent detection setting, intent labels are available for a subset of intents: there is an

additional labeled training set — which can be exploited, e.g., for (pre-)training a sentence encoder.

Zhang et al. (2022) evaluated their MTP and MTP-CLNN models by (pre-)training the encoder based on an unlabeled training set different from the test set where (new) intent clusters are induced, i.e., they evaluate on a held-out test set unseen during any (pre-)training phase. Since in our IDAS, no encoder is trained, we perform Steps 1–4 on the (unlabeled) test set following (Xie et al., 2016a; Yang et al., 2017a; Hadifar et al., 2019; Zhang et al., 2021a). To ensure a fair comparison we also consider an MTP-CLNN that uses that same test set in (pre-)training its encoder (i.e., for the $\mathcal{D}^{\text{unlabeled}}$ as defined in Zhang et al. (2022); results marked by ♠ in Table 3). Note that the test sets for a particular dataset are identical across all reported results.

First, we compare IDAS against the state-of-the-art in the *unsupervised* setting, i.e., MTP-CLNN, with results reported in Table 3. Both in the original settings of Zhang et al. (2022) (keeping the test data unseen during training, \diamond) as well as when using the unlabeled test data in training MTP(-CLNN) (♠), our IDAS significantly surpasses it, with gains averaged over three datasets of +3.19–3.94%, +1.79–2.86% and +1.96–3.34% in respectively ARI, NMI and ACC. We further find that IDAS consistently outperforms MTP-CLNN on all metrics and datasets, except for Banking, where IDAS and MTP-CLNN perform similarly (when comparing them in similar settings, i.e., both using unlabeled test data in training phase). Note that both IDAS and MTP-CLNN perform worse on StackOverflow and Banking in our settings (♠) compared to the original results of Zhang et al. (2022) (\diamond), likely because in case of ♠, the MTP(-CLNN) encoder(s) were trained on a substantially lower number of samples, i.e., only 5.5% for StackOverflow (1,000 for ♠ vs. 18,000 for \diamond) and 34% for Banking (3,080 for ♠ vs. 9,016 for \diamond).

Second, we assess our IDAS’s performance in the *semi-supervised* task setting, where a subset of intents has labeled data. Note however that our IDAS does not use the labels for those utterances in any way. The results for CLINC presented in Table 4 show that IDAS outperforms both semi-supervised SCL+PLT and DAC methods for KCR’s of 25% and 50%. Notably, IDAS surpasses SCL+PLT and DAC for KCR of 50%, with improvements in the range of 5.77–6.76%, 1.61–2.32%, and 4.78–4.89% in ARI, NMI, and ACC, respectively. Even for

Model	Banking			StackOverflow			Transport			Average		
	ARI	NMI	ACC	ARI	NMI	ACC	ARI	NMI	ACC	ARI	NMI	ACC
MTP \diamond	47.33	77.32	57.99	48.71	63.85	66.18	-	-	-	48.02	70.59	62.09
MTP-CLNN \diamond	<u>55.75</u>	<u>81.80</u>	<u>65.90</u>	<u>67.63</u>	<u>78.71</u>	<u>81.43</u>	-	-	-	<u>61.69</u>	<u>80.26</u>	<u>73.67</u>
IDAS	57.56	82.84	67.43	72.20	81.26	83.82	-	-	-	64.88 ± 1.07	82.05 ± 0.68	75.63 ± 0.82
MTP \spadesuit	39.52	72.03	51.66	29.66	47.46	48.97	44.51	74.71	57.51	37.90 ± 0.48	64.73 ± 0.31	52.69 ± 0.60
MTP-CLNN \spadesuit	<u>52.47</u>	<u>79.46</u>	64.06	<u>62.53</u>	<u>73.52</u>	<u>78.82</u>	<u>50.33</u>	<u>77.77</u>	<u>61.60</u>	<u>55.11</u> ± 1.32	<u>76.92</u> ± 0.74	<u>68.16</u> ± 1.02
IDAS	53.31	80.43	<u>63.77</u>	66.08	77.25	82.11	57.75	81.66	68.51	59.05 ± 1.92	79.78 ± 0.91	71.46 ± 1.57
Δ MTP-CLNN \diamond	+1.81	+1.04	+1.53	+4.57	+2.55	+2.39	-	-	-	+3.19	+1.79	+1.96
Δ MTP-CLNN \spadesuit	+0.84	+0.97	-0.29	+3.55	+3.73	+3.39	+7.42	+3.89	+6.91	+3.94	+2.86	+3.34

Table 3: Comparison against *unsupervised* state-of-the-art. \diamond : results from Zhang et al. (2022). \spadesuit : results from (pre-)training MTP(-CLNN) on the test set (rather than a distinct unlabeled training set). The **best** model is typeset in bold and the runner-up is underlined. Δ MTP-CLNN values are the absolute gains of our IDAS.

KCR	Model	CLINC		
		ARI	NMI	ACC
0%	SMPNET	63.82	89.01	71.30
	IDAS	79.02 ± 1.14	93.82 ± 0.38	85.48 ± 0.84
25%	DAC \heartsuit	65.36	89.12	75.20
	SCL+PLT \heartsuit	64.78	89.31	73.77
50%	DAC \heartsuit	72.26	91.50	80.70
	SCL+PLT \heartsuit	73.25	92.21	80.59
75%	DAC \heartsuit	79.56	93.92	86.40
	SCL+PLT \heartsuit	83.44	95.25	88.68

Table 4: Comparison against *semi-supervised methods* DAC and SCL+PLT. \heartsuit : results from Shen et al. (2021). Bold indicates **best** model. KCR: known class ratio.

KCR = 75%, it performs just slightly worse than DAC, further confirming IDAS’s effectiveness.

5.2 Ablations

Below, we investigate the impact of (i) the encoding strategies from §3.3, and (ii) ICL from §3.2 on IDAS’s performance. The results for each ablation are averaged over 5 runs with the utterances’ order corresponding to those used for presenting the main results, i.e., with IDAS’s default parameters values. Due to computation budget constraints, we only provide ablations on StackOverflow for (ii), since it requires GPT-3. For (i), we report results for Banking, StackOverflow, Transport, and CLINC.

Effect of the encoding strategies: Table 5 compares the cluster performance of these four encoding strategies: (1) $E(x)$ encodes only utterances; (2) $E(\ell)$ encodes only generated labels; (3) $\phi_{\text{AVG}}(x, \ell)$ (Eq. (1)) averages utterance and label encodings into a single vector representation;

(4) $\phi_{\text{SMOOTH}}(x, \ell)$ (Eq. (2)) smooths the averaged vector representations. All encoding methods leveraging the generated labels ℓ outperform the baseline $E(x)$ using only the utterance, leading to ARI, NMI, and ACC gains between 5.12–19.23%, 3.82–16.75%, and 4.32–13.87%, respectively. This confirms our main hypothesis that abstractly summarizing utterances improves intent discovery. Moreover, combining utterance and label encodings ($\phi_{\text{AVG}}(x, \ell)$) further improves upon using the label alone (performing on par only for CLINC). Adding smoothing ($\phi_{\text{SMOOTH}}(x, \ell)$) boosts performance even more.

Inferring the number of smoothing neighbors:

Smoothing requires selecting the number of neighbors n' . Our proposed IDAS selects the value of $n' \in \{5, \dots, 45\}$ that yields the highest silhouette score. To assess the effect of that chosen n' value, we plot the ARI, NMI, and ACC scores for varying n' in Fig. 2. We observe that the ARI, AMI, and ACC scores obtained with the automatically inferred n' are nearly identical to the best achievable performance, demonstrating that the silhouette score is an effective heuristic for selecting a suitable number of smoothing neighbors.

Random vs. nearest neighbor demonstrations:

IDAS employs KATE (Liu et al., 2022) to select the n ICL demonstrations most similar to x , i.e., $\mathcal{N}_n(x)$, for generating x ’s label (§3.2). To evaluate KATE’s effectiveness for intent discovery, we present results for IDAS where n (= 8) demonstrations are instead selected randomly. Table 6 shows a substantial improvement of KATE over the random selection method, where the latter only marginally outperforms IDAS *without* any demon-

Encoding	Banking			StackOverflow			Transport			CLINC		
	ARI	NMI	ACC	ARI	NMI	ACC	ARI	NMI	ACC	ARI	NMI	ACC
$E(x)$	47.33	77.32	57.99	48.71	63.85	66.18	44.51	74.71	57.51	63.82	89.01	71.30
$E(\ell)$	52.45	81.14	62.31	67.94	80.60	80.05	54.37	80.68	64.66	75.01	93.04	81.27
$\phi_{\text{AVG}}(x, \ell)$	54.47	82.35	63.25	69.20	80.76	81.29	55.91	81.11	65.94	75.65	93.33	81.04
$\phi_{\text{SMOOTH}}(x, \ell)$	57.56	82.84	67.43	72.20	81.26	83.82	57.75	81.66	68.51	79.02	93.82	85.48

Table 5: Effect of the encoding strategies.

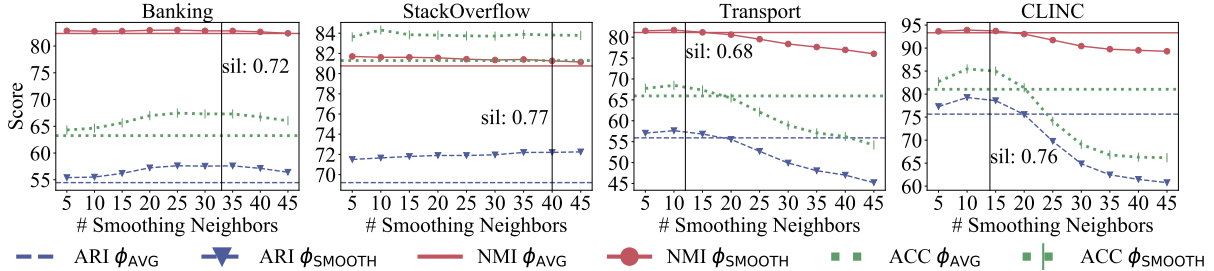


Fig. 2: Inferring the number of smoothing neighbors n' . The vertical lines represent the automatically determined number of smoothing neighbors corresponding to the highest silhouette score (sil).

strations (No ICL, $n = 0$). This follows the intuition that the LLM can pick a label from one of the n -NN instances, which likely shares an intent with the utterance to be labeled, thus effectively limiting label variation and improving clustering performance.

Varying the number of ICL demonstrations:

We generate labels (1) *without* ICL, adopting the static prompt for generating the prototypical labels, without any demonstrations, and (2) *with* ICL for varying numbers of demonstrations $n \in \{1, 2, \dots, 16\}$. Table 6 shows that (i) using any number of demonstrations leads to superior performance compared to using no demonstrations (No ICL); (ii) by varying small amounts of demonstrations ($n = 1, 2$, or 4) no significant differences are found; (iii) the best performance is achieved by using more demonstrations, i.e., 8 or 16. Consistent with the results of Min et al. (2022c); Lyu et al. (2022), increasing n from 8 to 16 does not result in further improvements, thus confirming that $n = 8$ demonstrations is a good default value.

Overestimating the number of prototypes: Following Hadifar et al. (2019); Zhang et al. (2021a,c, 2022), IDAS assumes a known number K of intents, both for the initial clustering (Step 1, retrieving prototypes, §3.1) and for the final clustering (Step 4, recovering latent intents, §3.4). While K can be *estimated* from a subset of utterances, determining it exactly is difficult. Unlike MTP-CLNN (Zhang

Method	StackOverflow		
	ARI	NMI	ACC
No ICL ($n = 0$)	66.21 \pm 0.13	77.27 \pm 0.04	80.42 \pm 0.13
KATE, $n = 1$	68.91 \pm 1.25	79.11 \pm 0.53	83.09 \pm 0.86
KATE, $n = 2$	68.88 \pm 1.40	79.06 \pm 0.86	82.67 \pm 0.98
KATE, $n = 4$	69.97 \pm 1.32	79.76 \pm 0.79	82.94 \pm 0.97
KATE, $n = 8$	<u>72.20\pm1.53</u>	<u>81.26\pm0.93</u>	<u>83.82\pm0.91</u>
KATE, $n = 16$	72.49 \pm 1.75	82.07 \pm 1.18	83.50 \pm 0.88
random, $n = 8$	66.80 \pm 0.90	78.72 \pm 0.85	81.37 \pm 0.93
$K \times 2$ ($n = 8$)	71.43 \pm 0.66	80.76 \pm 0.28	83.51 \pm 0.56

Table 6: ICL ablations. IDAS default settings are $n = 8$. The $K \times 2$ result uses twice the number of gold standard intents for the initial (Step 1, §3.1) clustering (i.e., 40 instead of 20 for StackOverflow).

et al., 2022), IDAS does not assume that the number of *samples* of each latent intent is known. To probe the robustness of IDAS’s label generation to an incorrect number of prototypes, we conduct the initial K -means clustering with twice the gold number of intents. The $K \times 2$ row in Table 6 shows that this results in only a minor performance drop, indicating that IDAS’s label generation process is sufficiently robust to such overestimation. In fact, we hypothesize that having multiple prototypes representing the same intent is less harmful than an insufficient number or incorrectly selected prototypes that do not accurately represent each intent.

6 Conclusions

Unlike existing methods that *train* unsupervised sentence encoders, our IDAS approach employs a *frozen* pre-trained encoder since it increases the (dis)similarity of (un)related utterances in the textual space by abstractly summarizing utterances into “labels”. Our experiments demonstrate that IDAS substantially outperforms the current state-of-the-art in unsupervised intent discovery across multiple datasets (i.e., Banking, StackOverflow, and our private Transport), and surpasses two recent semi-supervised methods on CLINC, despite not using any labeled intents at all. Our findings suggest that our alternative strategy of abstractly summarizing utterances (using a general purpose LLM) is more effective than the dominant paradigm of training unsupervised encoders (specifically on dialogue data), and thus may open up new perspectives for novel intent discovery methods. Since our generated labels provide a better measure of intent-relatedness, we hypothesize that they could also enhance the performance of existing methods that train unsupervised encoders, e.g., by (i) reducing the number of false positive contrastive pairs for MTP-CLNN (Zhang et al., 2022), or (ii) improving the purity of clusters induced by methods that iteratively cluster utterances and update the encoder with (self-)supervision from cluster assignments (Xie et al., 2016a; Caron et al., 2018b; Hadifar et al., 2019). To facilitate such follow-up work, we release our generated labels for the Banking, StackOverflow, and CLINC datasets.²

Limitations

Our work is limited in the following senses. First, all presented results relied on the ground truth number of intents to initialize the number of clusters for conducting K -means to retrieve prototypes (§3.1) and infer latent intents (§3.4). In practice, however, the ground truth number of intents is unknown and needs to be estimated by examining a subset of utterances. However, our ablations in §5.2 investigated the impact of overestimating the number of ground truth intents by a factor of two, and found that IDAS’s performance did not degrade much. While we did not explore this for the final K -means to infer latent intents, future work could investigate cluster algorithms that do not require the number of dialogue states as input, e.g.,

²<https://github.com/maarten-deraedt/IDAS-intent-discovery-with-abstract-summarization>.

DBSCAN (Ester et al., 1996), Mean shift (Comaniciu and Meer, 2002), or Affinity propagation (Frey and Dueck, 2007).

Second, we generated labels with the GPT-3 (175B) `text-davinci-003` model, which may be prohibitively expensive and slow to run for very large corpora. In our initial experiments, we tried using smaller-sized models such as `text-curie-001`, `text-babbage-001`, and `text-ada-001`, as well as Flan-T5-XL (Chung et al., 2022), but found that the generated labels were of lower quality compared to those of `text-davinci-003`. In future work, it would thus be interesting to further explore how to more effectively exploit such smaller-sized and/or open-source language models.

Ethics Statement

Since IDAS automatically recovers intents from utterances, e.g., those exchanged between users and support agents, any prejudices that may be present in these utterances may become apparent or even amplified in intents inferred by our model, since clearly IDAS does not eliminate such prejudices. Hence, when designing conversational systems based on such inferred intents, extra care should be taken to prevent them from carrying over to conversational systems deployed in the wild.

Moreover, since IDAS’s label generation process relies on LLMs, biases that exist in the data used to train these LLMs may be reinforced, leading to generated labels that may discriminate against or be harmful to certain demographics.

Acknowledgements

This work was funded in part by Flanders Innovation & Entrepreneurship (VLAIO), through Baeckeland project-HBC.2019.2221 in collaboration with Sinch Chatlayer; and in part by the Flemish government under the “Onderzoeksprogramma Artificiële Intelligentie (AI) Vlaanderen” program.

References

- Somnath Banerjee, Krishnan Ramanathan, and Ajay Gupta. 2007. Clustering short texts using wikipedia. In *Proceedings of the 30th annual international ACM SIGIR conference on Research and development in information retrieval*, pages 787–788.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *Advances in neural information processing systems*, 33:1877–1901.
- Mathilde Caron, Piotr Bojanowski, Armand Joulin, and Matthijs Douze. 2018a. Deep clustering for unsupervised learning of visual features. In *Proceedings of the European conference on computer vision (ECCV)*, pages 132–149.
- Mathilde Caron, Piotr Bojanowski, Armand Joulin, and Matthijs Douze. 2018b. Deep clustering for unsupervised learning of visual features. In *Computer Vision – ECCV 2018*, pages 139–156, Cham. Springer International Publishing.
- 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.
- Yanda Chen, Ruiqi Zhong, Sheng Zha, George Karypis, and He He. 2022. [Meta-learning via language model in-context tuning](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 719–730, Dublin, Ireland. Association for Computational Linguistics.
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Eric Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, Albert Webson, Shixiang Shane Gu, Zhuyun Dai, Mirac Suzgun, Xinyun Chen, Aakanksha Chowdhery, Sharan Narang, Gaurav Mishra, Adams Yu, Vincent Zhao, Yanping Huang, Andrew Dai, Hongkun Yu, Slav Petrov, Ed H. Chi, Jeff Dean, Jacob Devlin, Adam Roberts, Denny Zhou, Quoc V. Le, and Jason Wei. 2022. [Scaling instruction-finetuned language models](#).
- Dorin Comaniciu and Peter Meer. 2002. Mean shift: A robust approach toward feature space analysis. *IEEE Transactions on pattern analysis and machine intelligence*, 24(5):603–619.
- Alexis Conneau, Douwe Kiela, Holger Schwenk, Loïc Barrault, and Antoine Bordes. 2017. [Supervised learning of universal sentence representations from natural language inference data](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 670–680, Copenhagen, Denmark. Association for Computational Linguistics.
- Martin Ester, Hans-Peter Kriegel, Jörg Sander, and Xiaowei Xu. 1996. A density-based algorithm for discovering clusters in large spatial databases with noise. In *Proceedings of the Second International Conference on Knowledge Discovery and Data Mining*, pages 226–231.
- Brendan J Frey and Delbert Dueck. 2007. [Clustering by passing messages between data points](#). *science*, 315(5814):972–976.
- Tianyu Gao, Xingcheng Yao, and Danqi Chen. 2021. [SimCSE: Simple contrastive learning of sentence embeddings](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6894–6910, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Amir Hadifar, Lucas Sterckx, Thomas Demeester, and Chris Develder. 2019. [A self-training approach for short text clustering](#). In *Proceedings of the 4th Workshop on Representation Learning for NLP (RepL4NLP-2019)*, pages 194–199, Florence, Italy. Association for Computational Linguistics.
- Andreas Hotho, Steffen Staab, and Gerd Stumme. 2003. Ontologies improve text document clustering. In *Third IEEE international conference on data mining*, pages 541–544. IEEE.
- Xiaohua Hu, Xiaodan Zhang, Caimei Lu, Eun K Park, and Xiaohua Zhou. 2009. Exploiting wikipedia as external knowledge for document clustering. In *Proceedings of the 15th ACM SIGKDD international conference on Knowledge discovery and data mining*, pages 389–396.
- Ryan Kiros, Yukun Zhu, Russ R Salakhutdinov, Richard Zemel, Raquel Urtasun, Antonio Torralba, and Sanja Fidler. 2015. Skip-thought vectors. *Advances in neural information processing systems*, 28.
- Harold W Kuhn. 1955. The hungarian method for the assignment problem. *Naval research logistics quarterly*, 2(1-2):83–97.
- 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.
- Ting-En Lin, Hua Xu, and Hanlei Zhang. 2020. Discovering new intents via constrained deep adaptive clustering with cluster refinement. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 34, pages 8360–8367.

- Jiachang Liu, Dinghan Shen, Yizhe Zhang, Bill Dolan, Lawrence Carin, and Weizhu Chen. 2022. [What makes good in-context examples for GPT-3?](#) In *Proceedings of Deep Learning Inside Out (DeeLIO 2022): The 3rd Workshop on Knowledge Extraction and Integration for Deep Learning Architectures*, pages 100–114, Dublin, Ireland and Online. Association for Computational Linguistics.
- Stuart Lloyd. 1982. Least squares quantization in pcm. *IEEE transactions on information theory*, 28(2):129–137.
- Xinxi Lyu, Sewon Min, Iz Beltagy, Luke Zettlemoyer, and Hannaneh Hajishirzi. 2022. Z-icl: Zero-shot in-context learning with pseudo-demonstrations. *arXiv preprint arXiv:2212.09865*.
- George A Miller. 1995. Wordnet: a lexical database for english. *Communications of the ACM*, 38(11):39–41.
- Sewon Min, Mike Lewis, Hannaneh Hajishirzi, and Luke Zettlemoyer. 2022a. [Noisy channel language model prompting for few-shot text classification](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 5316–5330, Dublin, Ireland. Association for Computational Linguistics.
- Sewon Min, Mike Lewis, Luke Zettlemoyer, and Hannaneh Hajishirzi. 2022b. [MetalCL: Learning to learn in context](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2791–2809, Seattle, United States. Association for Computational Linguistics.
- Sewon Min, Xinxi Lyu, Ari Holtzman, Mikel Artetxe, Mike Lewis, Hannaneh Hajishirzi, and Luke Zettlemoyer. 2022c. [Rethinking the role of demonstrations: What makes in-context learning work?](#) In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 11048–11064, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Gray, et al. 2022. Training language models to follow instructions with human feedback. In *Advances in Neural Information Processing Systems*.
- Fabian Pedregosa, Gaël Varoquaux, Alexandre Gramfort, Vincent Michel, Bertrand Thirion, Olivier Grisel, Mathieu Blondel, Peter Prettenhofer, Ron Weiss, Vincent Dubourg, et al. 2011. Scikit-learn: Machine learning in python. *the Journal of machine Learning research*, 12:2825–2830.
- Nils Reimers and Iryna Gurevych. 2019. [Sentence-BERT: Sentence embeddings using Siamese BERT-networks](#). 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 3982–3992, Hong Kong, China. Association for Computational Linguistics.
- Peter J Rousseeuw. 1987. Silhouettes: a graphical aid to the interpretation and validation of cluster analysis. *Journal of computational and applied mathematics*, 20:53–65.
- Ohad Rubin, Jonathan Herzig, and Jonathan Berant. 2022. [Learning to retrieve prompts for in-context learning](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2655–2671, Seattle, United States. Association for Computational Linguistics.
- Xiang Shen, Yinge Sun, Yao Zhang, and Mani Nadjmabadi. 2021. [Semi-supervised intent discovery with contrastive learning](#). In *Proceedings of the 3rd Workshop on Natural Language Processing for Conversational AI*, pages 120–129, Online. Association for Computational Linguistics.
- Taylor Sorensen, Joshua Robinson, Christopher Rytting, Alexander Shaw, Kyle Rogers, Alexia Delorey, Mahmoud Khalil, Nancy Fulda, and David Wingate. 2022. [An information-theoretic approach to prompt engineering without ground truth labels](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 819–862, Dublin, Ireland. Association for Computational Linguistics.
- Karen Sparck Jones. 1972. A statistical interpretation of term specificity and its application in retrieval. *Journal of documentation*, 28(1):11–21.
- Douglas Steinley. 2004. [Properties of the hubert-adjustable adjusted rand index](#). *Psychological methods*, 9(3):386.
- Tingting Wei, Yonghe Lu, Huiyou Chang, Qiang Zhou, and Xianyu Bao. 2015. A semantic approach for text clustering using wordnet and lexical chains. *Expert Systems with Applications*, 42(4):2264–2275.
- Junyuan Xie, Ross Girshick, and Ali Farhadi. 2016a. [Unsupervised deep embedding for clustering analysis](#). In *International conference on machine learning*, pages 478–487. PMLR.
- Junyuan Xie, Ross Girshick, and Ali Farhadi. 2016b. [Unsupervised deep embedding for clustering analysis](#). In *Proceedings of The 33rd International Conference on Machine Learning*, volume 48 of *Proceedings of Machine Learning Research*, pages 478–487, New York, New York, USA. PMLR.
- Jiaming Xu, Peng Wang, Guanhua Tian, Bo Xu, Jun Zhao, Fangyuan Wang, and Hongwei Hao. 2015. [Short text clustering via convolutional neural networks](#). In *Proceedings of the 1st Workshop on Vector Space Modeling for Natural Language Processing*, pages 62–69, Denver, Colorado. Association for Computational Linguistics.

- Bo Yang, Xiao Fu, Nicholas D Sidiropoulos, and Mingyi Hong. 2017a. Towards k-means-friendly spaces: Simultaneous deep learning and clustering. In *international conference on machine learning*, pages 3861–3870. PMLR.
- Bo Yang, Xiao Fu, Nicholas D. Sidiropoulos, and Mingyi Hong. 2017b. [Towards k-means-friendly spaces: Simultaneous deep learning and clustering](#). In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 3861–3870. PMLR.
- Dejiao Zhang, Feng Nan, Xiaokai Wei, Shang-Wen Li, Henghui Zhu, Kathleen McKeown, Ramesh Nallapati, Andrew O. Arnold, and Bing Xiang. 2021a. [Supporting clustering with contrastive learning](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 5419–5430, Online. Association for Computational Linguistics.
- Hanlei Zhang, Xiaoteng Li, Hua Xu, Panpan Zhang, Kang Zhao, and Kai Gao. 2021b. [TEXTTOIR: 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*, pages 167–174, Online. Association for Computational Linguistics.
- Hanlei Zhang, Hua Xu, Ting-En Lin, and Rui Lyu. 2021c. [Discovering new intents with deep aligned clustering](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(16):14365–14373.
- Haode Zhang, Yuwei Zhang, Li-Ming Zhan, Jiaxin Chen, Guangyuan Shi, Xiao-Ming Wu, and Albert Y.S. Lam. 2021d. [Effectiveness of pre-training for few-shot intent classification](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 1114–1120, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Yuwei Zhang, Haode Zhang, Li-Ming Zhan, Xiao-Ming Wu, and Albert Lam. 2022. [New intent discovery with pre-training and contrastive learning](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 256–269, Dublin, Ireland. Association for Computational Linguistics.

Encoding	Banking			StackOverflow			CLINC			Average		
	ARI	NMI	ACC	ARI	NMI	ACC	ARI	NMI	ACC	ARI	NMI	ACC
MTP or paraphrase-mpnet-base-v2												
- $E(x)$	47.33	77.32	57.99	48.71	63.85	66.18	63.82	89.01	71.30	53.29	76.73	65.16
- $E(\ell)$	52.45	81.14	62.31	67.94	80.60	80.05	75.01	93.04	81.27	65.13	84.93	74.54
- $\phi_{\text{AVG}}(x, \ell)$	54.47	82.35	63.25	69.20	80.76	81.29	75.65	93.33	81.04	66.44	85.48	75.19
- $\phi_{\text{SMOOTH}}(x, \ell)$	57.56	82.84	67.43	72.20	81.26	83.82	79.02	93.82	85.48	69.59	85.97	78.91
all-mpnet-base-v2												
- $E(x)$	54.09	81.29	64.27	57.69	72.40	71.72	69.24	91.05	76.04	60.34	81.58	70.68
- $E(\ell)$	52.33	81.51	63.29	66.96	82.37	81.13	77.48	93.91	83.08	65.59	85.93	75.83
- $\phi_{\text{AVG}}(x, \ell)$	57.90	83.87	67.55	70.92	83.81	82.56	78.86	94.40	83.58	69.23	87.36	77.90
- $\phi_{\text{SMOOTH}}(x, \ell)$	59.88	84.13	70.07	78.27	85.09	87.02	82.26	94.93	87.80	73.47	88.05	81.84

Table 7: *Effect of using a more powerful sentence encoder.* The first four rows show the main results presented in §5.1, i.e., with the MTP encoder for Banking and StackOverflow, and with paraphrase-mpnet-base-v2 for CLINC. The last four rows show the results of performing the final clustering (Step 4) with encoder all-mpnet-base-v2.

A Appendix

In §A.1, we analyze how using a more powerful pre-trained sentence encoder affects the cluster performance of IDAS. Additionally, we present and discuss the prompts in §A.2, and conduct a qualitative analysis of the generated labels produced by our IDAS approach in §A.3. Finally, in §A.4, we provide a brief overview of the implementation details of our experiments.

A.1 Effect of using a more powerful encoder

Here, we assess the impact of using a more powerful frozen pre-trained encoder on the clustering performance of IDAS. Specifically, we provide results of the four encoding strategies using the SBERT encoder all-mpnet-base-v2 (Reimers and Gurevych, 2019) in Table 7. The overall results, presented in the three rightmost columns as the average of the scores across the three datasets, show that each encoding strategy for all-mpnet-base-v2 (bottom half of the table) consistently improves upon the corresponding results for the encoder used in our previous main results (as repeated here in the top rows). However, the label-only encoding strategy ($E(\ell)$) achieves similar results for different encoders, likely because the labels already are a short disambiguated version of their associated utterances. Conversely, the other three strategies that exploit the original utterances x deliver substantially better results for all-mpnet-base-v2, as the advanced encoder can more effectively disambiguate utterances based on their latent intents, thus improving cluster performance. Notably, using all-mpnet-base-v2 for the smoothing strategy ($\phi_{\text{SMOOTH}}(x, \ell)$) com-

pared to using MTP (Banking, Stackoverflow) or paraphrase-mpnet-base-v2 (CLINC), results in gains of +3.88%, +2.08%, and +2.93% in ARI, NMI, and ACC, respectively.

These results validate that employing more powerful pre-trained sentence encoders can further improve cluster performance out-of-the-box. It should be noted that, due to limitations in computation budget, we only replaced the encoder for Step 4 to induce intent clusters. However, we anticipate that using all-mpnet-base-v2 also for Steps 1–2 could result in additional improvements.

A.2 Prompts

Figures 3–4 present the static prompts used to generate prototypical labels in Step 2.1 (§3.2) without demonstrations, as well as the ICL prompts for generating labels of non-prototypical utterances in Step 2.2 (§3.2). One advantage of instructing LLMs is the ability to specify additional information in the prompts. When clustering topic datasets, there typically is a general understanding of the broad topic according to which utterances should be partitioned, and this topic can be specified in the prompts used to instruct LLMs. Since StackOverflow pertains to topics rather than intents, we adopted a more specific prototypical label generation prompt to instruct the LLMs to directly summarize the utterances based on the “technology” they refer to. While this approach may not be effective for intent discovery (i.e., a single conversational dataset can contain intents from multiple topics as well as non-topic intents), we speculate that it could be applied to other topic classification datasets, e.g., News or Biomedical, where a proto-

typical prompt could instruct the LLM to identify the “news category” or “medical drug”, “disease”, etc. We defer exploring IDAS for topic clustering beyond StackOverflow to future work.

A.3 Qualitative Analysis

We conduct a qualitative analysis of IDAS’s generated labels. Tables 8–10 show the generated labels for a subset of clusters induced in Step 4 for the corresponding StackOverflow, Banking, and CLINC datasets. For each presented cluster, we report (i) the generated labels with their associated counts in that cluster, and (ii) the majority gold intent, i.e., the most prevalent gold intent among utterances in that cluster, and the number of utterances within that cluster belonging to the majority gold intent.

Main findings: Overall, Tables 8–10 reveal that there is little variation among generated labels within a specific cluster. Specifically, for the majority of clusters, the most frequently occurring generated label has a notably higher count than the other generated labels, e.g., the first row in Table 8 shows that the label “*Magento*” is generated for 47 out of 49 utterances in that cluster. These findings further support our main hypothesis that abstract summarization increases the similarity in the input space of utterances with the same latent intent. Given the low variation across generated labels within clusters, we hypothesize that our generated labels could also make clusters more easy to interpret compared to utterance-only clustering, thereby potentially reducing the time required for manually inspecting clusters in real-world settings.

Slightly specific labels: While most clusters clearly contain a single label that appears much more frequently than other labels, there are some clusters, e.g., `pto_request`, `plug_type`, `reminder_update`, and `calories` for CLINC (Table 10), where this is not the case. However, a closer examination of these clusters reveals that the labels still exhibit low variation since they share the same syntactic and lexical structure. For instance, the `plug_type` cluster’s generated labels mostly follow the “Plug Converter ⟨noun adjunct⟩” pattern, with only the noun adjunct being specific to the utterance from which the label is generated. Note that for our intent discovery purpose, these slightly more specific labels do not negatively impact cluster performance, as long as there is a high overlap in syntactical and lexical structure among generated labels.

Overly general labels: Although some utterances are summarized into slightly more specific labels, others may be summarized into overly general labels. For instance, in the banking cluster `exchange_via_app` (Table 9) the label “*Foreign currency exchange*” appears 25 times. However, 6 of those 25 utterances do not have `exchange_via_app` as their gold intent, despite having obtained the same generated label as those other 19 utterances that do. This is due to the fact that generated labels corresponding to more high-level intents may be assigned to utterances that belong to different intents but share that common more high-level intent. For instance, the utterances “*Can this app help me exchange currencies?*” and “*I want to make a currency exchange to EU*” have respective gold intents `exchange_via_app` and `fiat_currency_support`, yet both are summarized into a more high-level “*Foreign currency exchange*” label. In contrast to generated labels that are slightly too specific, overly general labels can adversely affect cluster performance, as they may incorrectly group together utterances that belong to different intents despite sharing a common high-level intent.

A.4 Implementation Details

For all presented experiments, the utterances are encoded (Steps 1, 3–4) on a 2.6 GHz 6-Core Intel Core i7 CPU, using a frozen pre-trained sentence encoder. Similarly, both the initial and final K -means clustering to respectively retrieve prototypes (Step 1) and infer latent intents (Step 4), are conducted on CPU. We adopt the K -means implementation of `scikit-learn` (Pedregosa et al., 2011), with default parameter values, i.e., using the algorithm of Lloyd (1982) and `n_init = 10`.

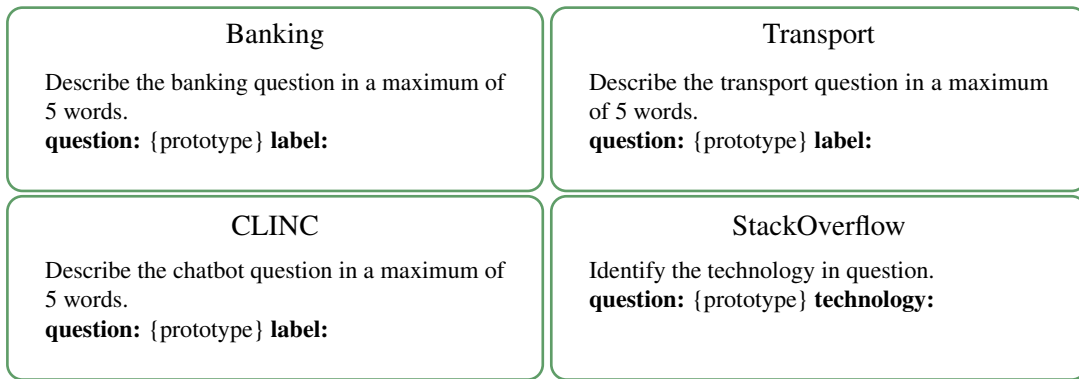


Fig. 3: *Static prototypical label generation prompts.* Note that since StackOverflow is a topic rather than an intent classification dataset, we adopt a slightly different prompt.

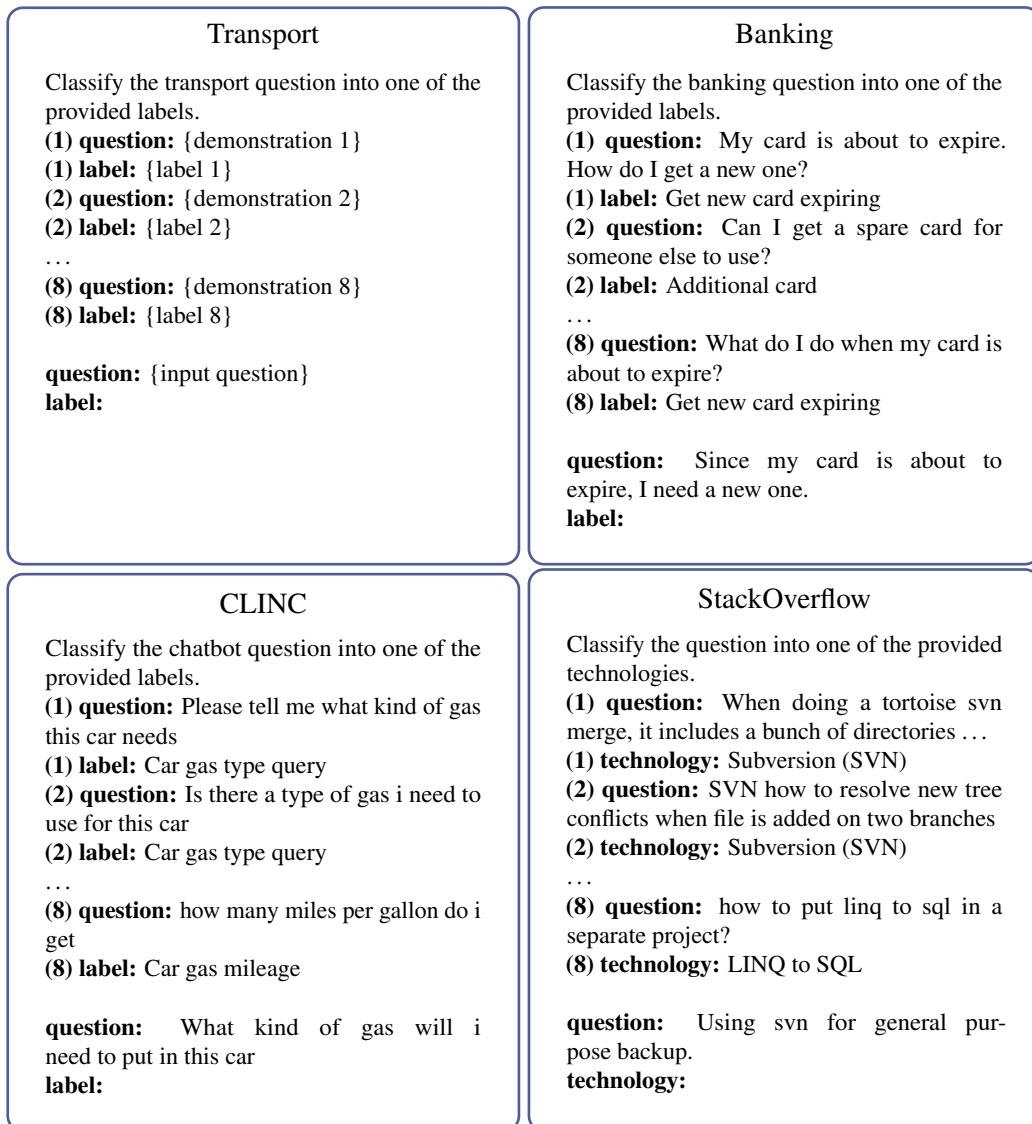


Fig. 4: *Prompts for non-prototypical label generation with ICL.*

Majority gold topic ($\#y_{\text{GOLD}}/ \mathcal{C} $)	Generated labels ($\#\ell$)		
topic_20 (49/49)	Magento (47)	Magento CodeIgniter (1)	Shipping Method (1)
topic_17 (44/45)	Drupal (35) Drupal and Ruby on Rails (1) Drupal and Microsoft SQL Server and Microsoft IIS 7 (1)	Drupal 6 (5) Drupal Ubercart (1)	Drupal 5 (1) Web View (1)
topic_10 (43/49)	BASH scripting (30) BASH scripting (2) Shell scripting (1) Scriptaculous (1) SSH scripting (1)	Shell Scripting (5) Scripting (1) Pipe-separated files (1) Shell Scripting (1)	Bash (Unix Shell) (2) Scripting (1) Readline (1) Bash scripting (1)
topic_6 (46/46)	Matlab (35) MATLAB (1) Matlab and C# (1) Ezplot (Matlab plotting tool) (1)	Matlab Octave (3) MatLab Mathematica (1) N/A (1)	Matrix (1) MatLab (1) Image Processing (1)
topic_19 (45/46)	Haskell (40) Haskell HDBC (1) GHCi (Glasgow Haskell Compiler Interactive) (1)	Haskell Cabal (1) General Programming (1) GHCi (Glasgow Haskell Compiler Interactive) (1)	General Programming (1)
topic_16 (42/45)	Qt (32) Qt4 (1) Qt (C++) (1) Real Time Video Capture (1)	Qt C++ (2) QT (1) QuickTime (1) QT (1)	Qt (C++ library) (2) QtScript (1) IP Camera (1) Quicksilver (1)
topic_1 (45/48)	WordPress (38) Open Atrium (1) HTTP POST (1) WordPress, RESTful, SOAP, InterWoven TeamSite (1)	jQuery and cycle (1) Disqus (1) Blogging (1) Commenting (1)	Drupal and WordPress (1) WordPress, PHP (1) WordPress and Django (1)
topic_5 (45/45)	Microsoft Excel (40) Microsoft Excel, Internet Information Services (IIS) (1)	Excel VBA (1) Microsoft Excel, Visual Basic (1)	Perl (1) Google Earth (1)
topic_3 (47/53)	Subversion (SVN) (42) Apache web server and Subversion (SVN) (1) Subversion (SVN) and WebDAV (1) Subversion (SVN) and Apache web server (1)	File System (3) Subversion (SVN) and SharpSvn (1) Subversion (SVN) and Windows (1) Concurrent Versions System (CVS) (1)	Subversion (SVN) (1) Version Control (1)

Table 8: Generated labels that occur in selected IDAS clusters for StackOverflow, as well as the number of times $\#\ell$ each label ℓ occurs in corresponding cluster \mathcal{C} . The majority gold topic y_{GOLD} of cluster \mathcal{C} is the most prevalent gold topic among all utterances in y_{GOLD} , and $\#y_{\text{GOLD}}$ denotes the number of utterances in \mathcal{C} with $y = y_{\text{GOLD}}$. Generated labels of utterances that have gold intents *different* than y_{GOLD} are highlighted in red. Since no descriptive topic names are provided for StackOverflow, we refer to them simply as numbered topics (topic_x)

Majority gold intent ($\# y_{\text{GOLD}}/ \mathcal{C} $)	Generated labels ($\# \ell$)	
lost_or_stolen_phone (38/38)	Lost phone banking app (37)	Switching phones banking app (1)
atm_support (35/35)	ATM card acceptance (25)	Find nearest ATM (10)
card_acceptance (24/27)	Card usage limits (24)	Card usage (3)
virtual_card_not_working (31/33)	Virtual card not working (31)	Virtual card not received (2)
contactless_not_working (37/39)	Contactless banking issue (37)	Banking login issues (2)
compromised_card (24/42)	Unauthorized card usage (24)	Unauthorized card usage (18)
age_limit (39/39)	Age requirement for banking (30)	Opening an account for family members (9)
terminate_account (40/41)	Close bank account (39) Change bank name (1)	Account closure advice (1)
card_about_to_expire (17/20)	Get new card expiring (17) Renew card banking (1)	Get new card swallowed (3)
card_delivery_estimate (13/13)	Delivery time in US (9) Delivery date selection (2)	Delivery time request (2)
country_support (17/17)	Banking countries operated in (14) Supported countries (1)	Banking locations (2)
automatic_topic (27/27)	Automated top-up option (14) Low balance top-up feature (5)	Auto top-up location query (7) Auto top-up activation issue (1)
receiving_money (14/18)	Banking - Salary Deposit (14) Banking - Types of Deposits (1)	Banking, Payment, Check (2) Banking, Deposit, Cheque (1)
receiving_money (10/19)	Configure salary in GBP (8) Convert currency to GBP (1) Convert currency to AUD (6)	Convert currency to GBP (2) Deposit Money in GBP (1) Convert currency to AUD GBP (1)
apple_pay_or_google_pay (40/40)	Top up with Google Pay (10) Apple Pay issue (10) Cost of Apple Pay (1)	Top up with Apple Pay (10) Top up with Apple Watch (8) Set up Apple Pay (1)
getting_spare_card (22/25)	Get second card banking (11) Link existing bank card (4) Get spare card banking (1)	Add card for family member (6) Link card to website (2) Choose bank card (1)
visa_or_mastercard (36/40)	Credit card offerings (19) Credit card application process (4) Credit card acceptance (1)	Credit card decision making (12) Card payment acceptance (3) Credit card eligibility (1)
balance_not_updated_after_cheque_or_cash_deposit (36/38)	Cash deposit not posted (25) Cheque deposit processing time (1) Cash deposit flagged (1) Direct Deposit not posted (1)	Cash deposit pending query (6) Cash deposit not accepted (1) Cash deposit to account (1)
exchange_via_app (27/51)	Foreign currency exchange (19) Currency conversion (1) Foreign currency exchange (6) Receive payment in foreign currency (5)	Currency exchange process (7) Cryptocurrency exchange (7) Cross-border payments (1) Discounts for frequent currency exchange (5)

Table 9: Generated labels that occur in selected IDAS clusters for Banking, as well as the number of times $\# \ell$ each label ℓ occurs in corresponding cluster \mathcal{C} . The majority gold intent y_{GOLD} of cluster \mathcal{C} is the most prevalent gold intent among all utterances in y_{GOLD} , and $\# y_{\text{GOLD}}$ denotes the number of utterances in \mathcal{C} with $y=y_{\text{GOLD}}$. Generated labels of utterances that have gold intents **different** than y_{GOLD} are highlighted in red.

Majority gold intent ($\# y_{\text{GOLD}}/ C $)	Generated labels ($\# \ell$)	
find_phone (15/15)	Locate Phone Request (15)	
vaccines (15/15)	Travel Vaccination Needed (15)	
exchange_rate (15/15)	Currency Exchange Rate (15)	
share_location (15/15)	Share Location Request (15)	
international_fees (15/15)	International Transaction Fees (15)	
report_fraud (13/13)	Fraudulent Transaction Inquiry (11)	Report Fraudulent Activity (2)
change_speed (15/15)	Speak slower please (8)	Speak faster please (7)
tire_pressure (15/15)	Tire Air Pressure Query (14)	Tire air pressure query (1)
international_visa (15/16)	Need International Visa (15)	Intercontinental Meaning (1)
pto_request_status (13/17)	Vacation Request Status (12) Vacation Request Process (3)	Vacation request status (1) Vacation Request (1)
weather (15/17)	Weather forecast query (14) AC Temperature Query (1)	Meteorological Data for Tallahassee (1) Set AC Temperature (1)
balance (14/15)	Bank Account Balance (11) bank account balance (1)	Check Account Balance (2) Bank Account Balance (1)
cancel_reservation (15/16)	Cancel restaurant reservation (8) Cancel Reservations (1) Cancel reservation for Network (1)	Cancel dinner reservation (4) Call restaurant to cancel reservation (1) Cancel Appointment (1)
pto_request (11/11)	PTO request for March (3) PTO request for June (2) PTO request for First to Ninth (1) PTO request for July (1)	PTO request for May (2) PTO request for January (1) PTO request for January to February (1)
plug_type (15/15)	Plug Type Query (3) Plug Converter El Salvador (1) Plug Converter Mexico (2) Plug Converter Denmark (1) Plug Converter Z (1) Plug Converter Guam (1)	Plug Converter Barcelona (2) Plug in electronics? (1) Plug Converter Thailand (1) Plug Converter Israel (1) Plug Converter Cairo (1)
reminder_update (14/28)	Ask Reminder List (9) Set Reminder (3) Confirm Reminder Laundry (1) Set Reminder Trash Out (1) Set Reminder Movie (1) Set Reminder Bring Jacket (1) Set Reminder Conference (1) Set Reminder Booking (1)	Remind of Forgotten Task (3) Set Reminder Later (2) Set Reminder Later (1) Set Reminder Dog Medicine (1) Set Reminder Pick Up Stan (1) Set Reminder Take Out Oven (1) Set Reminder Pay Bills (1)
calories (15/21)	Calorie content of apple (2) Calorie content of peanut butter (1) Calorie content of Coke (1) Calorie content of bacon (1) Calorie content of KitKat (1) Calorie content of Cheetos (1) Nutrition Info for Brownies (1) Health benefits of avocados (1) Health benefits of chocolate (1) Calorie content of Peanut Butter and Jelly Sandwich (1)	Caloric value of cookie (1) Calorie content of fries (1) Calorie content of whole cashews (1) Calorie content of cookie (1) Calorie content of bagels (1) Calorie content of chocolate ice cream (2) Nutrition Facts for Cheerios (1) Health benefits of apples (1) Nutrition Info for Lay's Potato Chips (1)

Table 10: Generated labels that occur in selected IDAS clusters for CLINC, as well as the number of times $\# \ell$ each label ℓ occurs in corresponding cluster C . The majority gold intent y_{GOLD} of cluster C is the most prevalent gold intent among all utterances in y_{GOLD} , and $\# y_{\text{GOLD}}$ denotes the number of utterances in C with $y=y_{\text{GOLD}}$. Generated labels of utterances that have gold intents **different** than y_{GOLD} are highlighted in red.