

Multi-Tenant Optimization For Few-Shot Task-Oriented FAQ Retrieval

Asha Vishwanathan
Verloop.io
asha@verloop.io

Rajeev Unnikrishnan Warriar
Verloop.io
rajeev@verloop.io

Gautham Vadakkekara Suresh*
Verloop.io
gautham.suresh09@yahoo.com

Chandra Shekhar Kandpal*
Verloop.io
chandrashekhar1503@gmail.com

Abstract

Business-specific Frequently Asked Questions (FAQ) retrieval in task-oriented dialog systems poses unique challenges vis-à-vis community based FAQs. Each FAQ question represents an intent which is usually an umbrella term for many related user queries. We evaluate performance for such Business FAQs both with standard FAQ retrieval techniques using query-Question (q-Q) similarity and few-shot intent detection techniques. Implementing a real-world solution for FAQ retrieval in order to support multiple tenants (FAQ sets) entails optimizing speed, accuracy and cost. We propose a novel approach to scale multi-tenant FAQ applications in real-world context by contrastive fine-tuning of the last layer in sentence Bi-Encoders along with tenant-specific weight switching¹.

1 Introduction

Business-specific Frequently Asked Questions (FAQ) form an important part of many task-oriented dialog systems today. Business FAQs exhibit some common characteristics which differentiate them from community-based FAQs. We show few examples of such FAQs in Table 1. The intent is often a user-defined umbrella term that represents a group of questions in such FAQs. The system needs to respond to user queries which indicate the same intent as that of the FAQ question (Example-1). In some cases, similar FAQ questions can lead to different answers (Example-2). The answers may not always have overlap with the question (Example-3). The response can also be modeled as a system action rather than a text response (Example-4). The system should also be able to identify Out-Of-Scope (OOS) queries and redirect to a human agent.

Since each question within an intent category represents it loosely, having a single question for

an intent poses a challenge for intent detection. Simple paraphrases do not suffice and therefore, 5-10 variations for each FAQ question representing the intent is sourced from domain experts.

Common approaches in FAQ retrieval include Query-Question similarity and Query-Answer entailment (Sakata et al., 2019; Hammond et al., 1995; Tomuro and Lytinen, 2004). Query-Question (q-Q) methods focus on similarity between a user’s query and a question to retrieve the relevant answer. Query-Answer (q-A) methods predict the relevance between the query and an FAQ answer to re-rank the results obtained via q-Q. For FAQ retrieval, we leverage q-Q similarity. The q-A entailment method is infeasible in some FAQ sets due to the generic nature of answers for several different questions, owing to little match between the q-A pair and responses being modeled as system action instead of text.

We also prefer a single model based approach over stacked models such as retriever-reranker models (Zhang et al., 2020; Gupta and Carvalho, 2019) as inference latency gets compounded in such systems. We compare several q-Q retrieval approaches using pre-trained models and fine-tuned embeddings. We also model the problem as a few-shot intent detection and evaluate classifier based techniques.

From an industry specific dialog solutions perspective, the approaches need to balance accuracy, real time inference latency and costs while having the ability to scale for multiple tenant (FAQ sets) requirements. Our contribution is towards implementing an FAQ retrieval system for Business FAQs in real-world contexts. We describe a low-cost approach to deploy multi-tenant FAQ applications at scale while optimizing for inference latency and accuracy. We also evaluate different methods and models against few-shot intent detection and conversational datasets.

*Work done while the authors were working at Verloop.io

¹<https://github.com/verloop/few-shot-faqir>

No.	Question	Answer
1	How to check exchange rate? Rate for Peso	Please use the Rate Enquiry option on the Main Menu
2	Am I going to get a refund When am I going to get the refund	We issue refund for specific products Refund normally takes 7 days to process
3	My transaction failed My checkout did not happen properly	Please contact on XXX or mail us at XXX
4	Can I speak to an agent?	System action: Transfer to agent

Table 1: Examples of actual business questions. All identifying information is masked.

1.1 Real world constraints

We process millions of chat messages every day, of which a sizable percentage is related to FAQ retrieval. The expected latency requirement is less than 100 ms under production loads. Tenant specific FAQ sets vary in terms of the domain and intent granularity. It is important to maximize the retrieval accuracy for each tenant while not increasing inference costs drastically.

As on-premise machines are expensive to set-up and maintain, it is often cost-effective and convenient to leverage cloud providers. The application efficiently needs to utilize machine resources and scale for multiple tenants without significant overhead. Since GPU costs are higher, it is cost-effective to train using GPU machines and perform inference with CPU machines.

Apart from this, many tenants require the system to respond with certain standardized answers based on company policy without too many modifications. Therefore, retrieval based solutions are more acceptable than generative solutions. The problems of hallucinations which still exist in generative solutions (Roller et al., 2021; Shuster et al., 2021) constraints the use of these approaches for task-oriented FAQ responses in business contexts.

2 Related Work

The models based on transformer architectures have enabled the usage of transfer learning for various tasks such as text similarity (Devlin et al., 2019). Sakata et al. (2019) introduced an approach to combine the query-question similarity from TSUBAKI and query-answer relevance from BERT. Arora et al. (2020) evaluated different providers on three intent recognition datasets, created using queries received by chatbots from real users.

In a few-shot setting, recent approaches leverage the transformer architecture. Casanueva et al. (2020) demonstrated that using a dual sentence en-

coder as base model works better in low-resource intent detection tasks as compared to BERT. Zhang et al. (2021b) proposed the use of contrastive learning to improve the performance of BERT based classifiers in few-shot intent detection tasks.

Pre-training methods for creating transferable language representations is a common approach. In-domain pre-training further increases the adaptability of the model to domain related downstream tasks. Specifically for few-shot setting, fine-tuning pre-trained models with supervised learning is found to be effective (Zhang et al., 2021a).

Cross-encoder approaches with BERT (Devlin et al., 2019) and RoBERTa (Liu et al., 2019) models have achieved SOTA scores using sentence pair training. Training is done with question pairs. During inference, question pairs are created using query and train questions. These pairs are scored by a classifier to predict the final label.

Bi-Encoder models have the advantage of being able to cache the representations and hence, are highly efficient during inference (Reimers and Gurevych, 2019). The research on STILTS (Phang et al., 2018) found the advantages of intermediate training on data rich supervised tasks, giving pronounced benefits in data constrained regimes. Casanueva et al. (2020) discusses the gains achieved via transfer learning using pre-trained sentence encoders as a retriever model. Cross-Encoder models entail a higher latency as all the question pairs are required to be classified during inference. Zhang et al. (2020) explores usage of the classifier model as a re-ranking model following a Bi-Encoder based retriever to reduce inference time.

We evaluate our FAQ retrieval use case with simple baseline approaches, in-domain pre-trained models, classifier based approaches, Cross-Encoders and Bi-Encoders in a few-shot setting.

3 Datasets

We choose datasets that are similar to our real-world tenant FAQs and perform experiments on publicly available intent detection datasets for task-oriented dialogues. We use datasets from HINT3 (Curekart, Powerplay11 and SOFMattress) (Arora et al., 2020) which reflects real-world FAQ intents and user queries. We also use DialoGLUE datasets (Mehri et al., 2020) - HWU64 (Liu et al., 2021), CLINC150 (Larson et al., 2019) and BANKING77 (Casanueva et al., 2020).

DialoGLUE datasets further provide train sets for data constrained regimes with 5 and 10 samples per class. Similarly, HINT3 also contains subset variants of datasets which are created by retaining only samples whose entailment score using the ELMo model was less than 0.6. We use the few-shot subsets for DialoGLUE and the full as well as subset variant of HINT3 datasets. These datasets are suitable for few-shot analysis. All the datasets have separate train and test sets and we report our evaluations on the test sets. The CLINC150 and HINT3 datasets have separate in-scope and out-of-scope queries and we use them to evaluate OOS performance. The dataset size, intent and domain distribution is listed in Appendix A.

4 Methodology

In this section, we elaborate on our experiments for model training and deployment. We compare different approaches of few-shot FAQ retrieval and describe our deployment strategy for the same.

4.1 FAQ Retrieval approaches

4.1.1 Baseline

We consider BM25 (Mass et al., 2020; Karan and Šnajder, 2016) and vector space using TF-IDF (Karan and Šnajder, 2016) as baseline approaches.

4.1.2 Pre-trained and fine-tuned features

The following models are considered for feature extraction:

- Pre-trained BERT
- Pre-trained ConvBERT: ConvBERT model is pre-trained on open-domain conversational data (Mehri et al., 2020).
- Fine-tuned ConvBERT: We further fine-tune ConvBERT model as an intent classifier and use the fine-tuned encoder as a feature extractor.

- Pre-trained Bi-Encoders: We use pre-trained Sentence BERT models (*mini-LM-L6-v2*, *all-mpnet-base-v2*) which are trained on more than 1 Billion sentence pairs consisting of a diverse set of duplicate question pairs, NLI sets, QA pairs.
- Fine-tuned Bi-Encoders: We fine-tune pre-trained Bi-Encoders for a Semantic textual similarity task with a contrastive loss (Zhang et al., 2021b).
- Fine-tuned Task adapted pre-trained Bi-Encoders: We perform pre-training for adaptation of pre-trained Bi-Encoders with a triplet loss using a cosine distance metric and further fine-tune using a contrastive loss

After the features are extracted, inference is implemented using a cosine similarity of query and question embeddings.

4.1.3 Classifier

We fine tune BERT with a linear layer on top to predict the intent class directly (Mehri et al., 2020).

4.1.4 Cross-Encoders

Inspired by the success of Cross-Encoders and the STILTS approach, we fine-tune a pre-trained Cross-Encoder model from SBERT (*stsb-distilroberta-base*).

4.2 Question Pair/Triplet Based fine-tuning

We form Question pairs to fine-tune all Bi-Encoder and Cross-Encoder models. We mark the question pairs belonging to same class as positive samples and ones belonging to different classes as negative samples.

If $c \in C$ Where C is the set of all classes, then we denote $q_{c,i}$ to be the i^{th} element of the c^{th} class. Then, $q_{c,i} \in Q$ where Q is the set of all Questions. $i \in \{1, \dots, n\}$ where $n \in \{5, 10, N\}$ when testing for few-shot training with 5 samples, 10 samples or the entire set respectively. The question pairs are labelled as defined below.

$$L(q_{l,j}, q_{m,k}) = \begin{cases} 1 & \text{if } l = m \\ 0 & \text{if } l \neq m \end{cases} \quad (1)$$

$$\forall l \neq m \text{ and } j \neq k$$

This method increases the data size as it generates ${}^n C_2$ question pairs given n questions. Even in a few-shot setting where there are only 180 samples across 21 intents (SOFMattress), 16110 question

pairs were generated. The Q-Q data can grow exponentially, so the train data is capped to 200K for fine-tuning. We also check fine-tuning on 50K Balanced samples. Hard sampling is done using a probabilistic method based on sample weights. For Q-Q pair with label 0, the weight is equal to the cosine similarity between the Q-Q pair. For label 1, the weight is (1 - cosine similarity) of the Q-Q pair. We use the same pre-trained bi-encoder model for the embeddings which we intend to fine-tune. Sampling with replacement is then done for labels 0 and 1 to get the required sample sizes. For the triplet based pre-training, the triplets are constructed such that the anchor and positive belong to the same label and the negative belong to a different label. Q-Q pairs are first constructed using the hard sampling approach described above. For each sample which is the triplet anchor, we get its positives and negatives from the Q-Q pairs. Sampling is then done based on weightage to construct the triplets.

4.3 Training and deployment

The common approach for fine-tuning models involves modifying weights in all layers of the pre-trained model. Instead, we propose an approach where only the final layer of the model is fine-tuned. The weights of all layers which are not fine-tuned are shared across tenant models. During inference in the production environment, we load the base model only once for shared parameters. We keep the tenant-specific weights for all tenants loaded in memory and we switch the tenant specific weights in the model for every inference request. The shared base model weights reduce the memory requirement by a significant margin when we scale to a large number of tenants. Since we retain the tenant weights in memory, replacing the model weights does not result in any significant latency overhead. This allows us to support multiple tenants using the same machine and in-turn reduces the inference costs by an order based on the model size. We select the best model based on expected number of tenants, throughput and memory requirement. We evaluate the performance of all fine-tuned approaches under the following constraints: Freezing all layers of the encoder except the last layer, training for a fixed number of iterations and using few-shot samples.

5 Experimentation Approaches

Fine-tuning for all Bi-Encoder and Cross-Encoder models is done using question pairs as elaborated in Section 4. Based on the training method outlined by Mosbach et al. (2021), we train for a higher number of iterations to offset any fine-tuning instability. For the Bi-Encoders, we fine-tune only the final layer of the model using a contrastive or an online triplet loss. The pre-trained BERT and ConvBERT classifier models were fine-tuned using a softmax cross-entropy loss. The Cross-Encoder was also fine-tuned with the same approach using question pairs with a binary cross-entropy loss. For predicting the query label, we create q-Q pairs limited to 5 questions per intent. The model predicts a score for these combinations, where the label for the question with the highest score is considered to be the predicted label. We trained all models for 10 K iterations with a learning rate of $2e-5$, batch size of 16, AdamW optimizer with 10% linear warm up and gradient normalization.

5.1 Pre-training

Commonly pre-training is done for in-domain adaptation using unlabelled datasets. But in our case, tenant FAQs are spread across multiple niche domains, making it difficult to get domain related data. Moreover, pre-training for each domain separately would result in multiple models per domain, making the hosting costs higher. In case of pre-trained dense retrievers, training on the base language model in an MLM fashion requires a further training of the bi-encoder model. In the GPL paper, Wang et al. (2021) adopt a pre-training method which is a triplet based training where the triplets are constructed from the unlabelled data. Gururangan et al. (2020) describe an approach of a second level pre-training using unlabelled corpus and pre-training using task related samples within a domain.

We experiment with a similar approach as GPL but we use the labelled data for the same task available across different domains. We construct triplets using an offline approach instead of an online batch mode, to ensure in-domain triplets. We create 100 K triplets from each dataset and get a total training corpus of 600 K samples. We do a second level of pre-training of the pre-trained bi-encoder models for 140 K iterations, with the Triplet loss using Cosine distance metric and a margin of 0.15. We then use this pre-trained model for further finetuning

Method	BANKING77		CLINC150		HWU64	
	5	10	5	10	5	10
BM25	53.96	61.10	55.37	60.80	50.37	54.46
TF-IDF	49.51	58.14	60.91	58.14	60.91	54.36
BERT - MP	40.25	48.89	68.97	66.73	52.88	57.06
ConvBERT - MP	50.42	59.41	73.11	79.53	58.92	65.79
ConvBERT - FT(C)	56.98	66.91	76.44	82.91	64.86	71.46
SBERT (MiniLM-L6)	76.78	83.47	79.08	81.24	68.02	73.42
SBERT (MiniLM-L6) - FT(C)	80.74	86.00	84.55	87.13	76.20	79.73
SBERT (MiniLM-L6) - FT(C) 50K	76.42	83.18	83.97	86.86	73.97	77.60
SBERT (MiniLM-L6) - FT(T)	81.33	86.33	85.11	87.75	75.27	82.24
SBERT (MiniLM-L6) - PT-FT(C)	84.28	84.67	89.88	89.86	85.68	86.24
SBERT (MPNet) - FT(T)	83.21	88.18	88.68	91.00	78.06	83.05
SBERT (MPNet) - PT-FT(C)	86.98	87.27	92.51	92.68	86.24	85.5
BERT - FT(C)	22.13	23.64	17.35	15.12	39.98	41.13
SBERT Cross-Encoder - FT(C)	67.10	69.83	76.10	75.20	66.91	68.03

Table 2: Top-1 accuracy of models on the DialoGLUE test sets. MP stands for Mean-Pooling, FT for fine-tuning, (C) for Contrastive loss, (T) for Triplet loss and PT for pre-training. Here, 5 and 10 refers to training subsets created with 5 and 10 samples per intent, respectively.

with a contrastive loss.

6 Evaluation Metrics

For evaluation, we consider a combination of FAQ Retrieval metrics along with Intent detection metrics. Similar to the metrics used in Sakata et al. (2019), we consider Top-k accuracy (same as Success Rate referred in the paper), MRR@k (Mean Reciprocal Rank), nDCG@k (normalized Discounted Cumulative Gain) and MAP@k (Mean Average Precision). From an intent detection perspective, we also evaluate out-of-scope accuracy across different thresholds. We use accuracy as the metric for evaluation of OOS queries. In case the top predicted label score is less than a specified threshold, then the top-k results are shown as suggestions to the user. Metrics like MRR, nDCG and MAP put emphasis on the ordering of the final responses and hence are useful for evaluating the effectiveness of this approach.

7 Results

Table 2 and 3 show the Top-1 accuracy for all experiments. Table 4 shows the performance of the MiniLM-L6 model in top-3 setting. Figure 1 further shows its OOS and in-scope accuracy against different thresholds.

From our experiments, we find that the fine-tuned Bi-Encoder models perform the best across all datasets. All the Bi-Encoder embedding ap-

proaches significantly outperform the other approaches in few-shot setting in spite of being constrained with the last layer fine-tuning strategy.

For HINT3 datasets, we use the benchmarking results as reported by Arora et al. (2020) for comparison. We find that the fine-tuned Bi-Encoder shows better performance than all benchmarked chatbot solutions on both full and constrained datasets. We noticed improvement while fine-tuning the pre-trained Bi-Encoder models using the Question pair approach. We also note that online triplet loss based fine-tuning performed marginally better as compared to contrastive Loss for all the base models other than MPNet model on Haptik dataset. We observe that the question pair / triplet training strategy mitigates the effect of fewer samples. Reducing sample sizes of the Q-Q pairs from 200 K to 50 K balanced samples hurts the performance across all datasets.

Fine-tuning the task adapted pre-trained models shows much better results in comparison to the base pre-trained models in most of the Dialogue datasets but the gains are either less or deteriorate for Haptik datasets. In Haptik datasets, the median examples per intent is quite low as compared to Dialogue, with some intents having only 1 sample. There would be no triplets formed for such samples. We also see a wide variation in the number of examples per intent ranging from 3 to 95 in the Curekart dataset. This would reduce the number of

Method	Curekart		Powerplay11		SOFMattress	
	Full Set	Subset	Full Set	Subset	Full Set	Subset
BM25	72.34	71.20	51.63	49.09	58.44	52.24
TF-IDF	72.56	70.35	53.81	52.72	59.74	54.54
BERT - MP	52.87	50.21	30.54	26.18	38.09	32.03
ConvBERT - MP	65.92	63.27	42.18	37.45	49.35	47.18
ConvBERT - FT(C)	77.40	75.22	46.9	43.6	62.33	56.27
SBERT (MiniLM-L6)	82.52	79.64	64.00	62.54	74.58	71.42
SBERT (MiniLM-L6) - FT(C)	87.38	86.06	64.00	63.2	78.78	77.48
SBERT (MiniLM-L6) - FT(C) 50K	86.28	86.94	61.45	58.9	75.32	74.89
SBERT (MiniLM-L6) - FT(T)	86.50	85.17	61.80	64.36	77.90	74.02
SBERT (MiniLM-L6) - PT-FT(C)	85.39	85.39	62.54	61.81	73.16	73.59
SBERT (MPNet) - FT(T)	85.61	84.51	66.54	64.00	74.45	73.59
SBERT (MPNet) - PT-FT(C)	88.05	87.83	62.18	61.81	75.32	76.19
BERT - FT(C)	55.75	58.63	18.18	18.91	41.56	38.96
SBERT Cross-Encoder - FT(C)	65.48	67.92	51.27	49.09	58.40	60.60
Dialogflow	75.00	71.20	59.60	55.60	73.10	65.30
Rasa	84.00	80.50	49.00	38.50	69.20	56.20
LUIS	59.30	49.30	48.00	44.00	59.30	49.30
Haptik	72.20	64.00	66.50	59.20	72.20	64.00
BERT Full Training	73.50	57.50	58.50	53.00	73.50	57.10

Table 3: Top-1 accuracy of models on HINT3(v1) test sets with a threshold of 0.1. Here, MP stands for Mean-Pooling, FT for fine-tuning, (C) for Contrastive loss, (T) for Triplet loss and PT for pre-training.

triplets for such intents.

While Cross-Encoder strategies are supposed to be superior, we find that the Cross-Encoder fine-tuning suffers due to the restrictive training strategy and performs sub-par when compared to the Bi-Encoders.

The BERT based classifier also shows poor performance which appears to degrade on datasets with more classes. In comparison to training BERT model with all layers unfreezed as reported in the [Arora et al. \(2020\)](#), we see that the last layer fine-tuning strategy severely impacts the model performance.

We also observe that models pre-trained on in-domain data such as ConvBERT is superior to base BERT even as a feature extractor. Interestingly, we see that even baseline approaches with BM25 and TF-IDF show better results than pre-trained BERT in such settings. Fine-tuning the ConvBERT model on the supervised classification task shows an improvement over the pre-trained model. If we look at the OOS accuracy for the fine-tuned MiniLM-L6 model, the similarity scores are normally high and lower thresholds do not have much impact. We further observe that each dataset has different thresholds where the best trade-off between OOS

and in-scope accuracy is achieved. The top-3 accuracy of MiniLM-L6 on all datasets including the challenging Powerplay11 dataset is above 70%. We choose the fine-tuned Bi-Encoder models based on superior performance with our training strategy.

8 Deployment

Smaller models give us more scalability at the cost of accuracy while larger models tend to have lower throughput and higher memory requirements. The model selection is done based on the specific business need which dictate these three metrics. Table 5 shows the load testing results using Locust framework, where we measure the mean latency per request and the request per second (RPS) achieved. The models were tested under increasing loads until RPS was stagnant and latency shot up. The results depict the performance at the maximum load the model can comfortably handle. The models are hosted via Kubernetes and deployed as a microservice using native Pytorch inference. Machine configuration used for locust testing was c2-standard-4 machine (4 vCPU and 16GB RAM) on Google Cloud Platform. From Table 6, we find that weight switching saves 81.5% memory per tenant for MiniLM-L6. For the larger model MPNet,

Metrics	Curekart		Powerplay11		SOFMattress		CLINC150	
	Full	Subset	Full	Subset	Full	Subset	Sample-5	Sample-10
Success Rate	89.80	88.93	73.81	73.09	81.38	81.81	94.22	95.08
MRR	88.38	87.38	68.42	67.45	79.94	79.43	88.88	90.77
nDCG	89.33	88.69	72.74	71.48	81.06	81.33	93.09	94.33
MAP	88	87	68	67	80	79	89	91

Table 4: Top-3 accuracy of fine-tuned MiniLM-L6 on the OOS datasets




Figure 1: In-scope and OOS accuracy on different datasets for fine-tuned MiniLM-L6 model.

it saves 92.5% per tenant. The benefits of weight switching increases as model sizes gets larger. We see that fine-tuned Bi-Encoder models work best in terms of accuracy in few-shot setting. We also have a latency benefit as the question embeddings are computed upfront and cached. At inference, prediction are done as cosine similarity between the query embeddings and the cached question embeddings. For further gains, we cache these tenant specific embeddings using FAISS or ANNOY². While the task adapted pre-trained models worked quite well in terms of accuracy, it entails repeated pre-training of the model for every new tenant from a niche domain and the subsequent fine-tuning. It poses a hindrance in decoupling tenant data seamlessly. However, it is still a very viable approach in case of separate hosting of domain specific models.

Model	Median latency	90%ile latency	RPS
MiniLM-L6	93	260	22.9
MPNet	210	500	6.2

Table 5: Load testing results with latency and requests per second (RPS) for best models. All latency values in milliseconds (ms).

Model	Memory requirement	
	Full model	Weight switch
MiniLM-L6	110	20.5
MPNet	560	41.5

Table 6: Memory increase per tenant for the best performing models. All values are in MegaBytes (MBs)

We choose MiniLM-L6 as it handles higher load at significantly lower latency as compared to the best performing model MPNet, with a slight trade-off in accuracy.

9 Conclusion

We evaluated various methods for retrieval of Business FAQs by modeling the problem as an FAQ retrieval and a few-shot intent detection task. We proposed a realistic multi-tenant deployment solution with trade-offs in accuracy while balancing for cost and latency. Our last layer weight-switching strategy works well where the model has to be fine-tuned on tenant specific tasks. We observed that fine-tuned Bi-Encoder embeddings work best in a few-shot setting even under a constrained fine-tuning strategy.

²<https://github.com/spotify/annoy>

10 Limitations

Although the weight switching approach is highly scalable with increase in number of tenants, it has few limitations. Our last-layer weight switching strategy is more effective for heavier models and the benefit starts to diminish with lighter models. Although this approach is highly scalable, implementing it for inference frameworks such as ONNX Runtime and NVIDIA Triton is not straightforward. Hence, we are currently limited to using native PyTorch inference which has higher latency. Apart from this, fine-tuning only the last layer also constraints the model training leading to lower accuracy as compared to training all the layers. All our experiments showcase benefits only with base models and do not showcase inference benefits using other strategies such as quantization, pruning, ANNs and caching. Such strategies can make the system more scalable.

11 Acknowledgements

This work was supported by Verloop.io. We wish to thank Gaurav Singh and Peeyush Jain from Verloop. We thank the anonymous reviewers whose thoughtful comments improved our final work.

References

- Gaurav Arora, Chirag Jain, Manas Chaturvedi, and Krupal Modi. 2020. [HINT3: Raising the bar for intent detection in the wild](#). In *Proceedings of the First Workshop on Insights from Negative Results in NLP*, pages 100–105, Online. Association for Computational Linguistics.
- 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.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Sparsh Gupta and Vitor R. Carvalho. 2019. [Faq retrieval using attentive matching](#). In *Proceedings of the 42nd International ACM SIGIR Conference on Research and Development in Information Retrieval, SIGIR'19*, page 929–932, New York, NY, USA. Association for Computing Machinery.
- Suchin Gururangan, Ana Marasović, Swabha Swayamdipta, Kyle Lo, Iz Beltagy, Doug Downey, and Noah A. Smith. 2020. [Don't stop pretraining: Adapt language models to domains and tasks](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8342–8360, Online. Association for Computational Linguistics.
- K. Hammond, R. Burke, C. Martin, and S. Lytinen. 1995. [Faq finder: a case-based approach to knowledge navigation](#). In *Proceedings the 11th Conference on Artificial Intelligence for Applications*, pages 80–86.
- Mladen Karan and Jan Šnajder. 2016. [Faqir – a frequently asked questions retrieval test collection](#). In *Text, Speech, and Dialogue*, pages 74–81, Cham. Springer International Publishing.
- 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.
- Xingkun Liu, Arash Eshghi, Pawel Swietojanski, and Verena Rieser. 2021. [Benchmarking Natural Language Understanding Services for Building Conversational Agents](#), pages 165–183. Springer Singapore, Singapore.
- Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike Lewis, Luke Zettlemoyer, and Veselin Stoyanov. 2019. [Roberta: A robustly optimized bert pretraining approach](#).
- Yosi Mass, Boaz Carmeli, Haggai Roitman, and David Konopnicki. 2020. [Unsupervised FAQ retrieval with question generation and BERT](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 807–812, Online. Association for Computational Linguistics.
- S. Mehri, M. Eric, and D. Hakkani-Tur. 2020. [Dialogue: A natural language understanding benchmark for task-oriented dialogue](#). *ArXiv*, abs/2009.13570.
- Marius Mosbach, Maksym Andriushchenko, and Dietrich Klakow. 2021. [On the stability of fine-tuning {bert}: Misconceptions, explanations, and strong baselines](#). In *International Conference on Learning Representations*.

- Jason Phang, Thibault F3evry, and Samuel R. Bowman. 2018. [Sentence encoders on stilts: Supplementary training on intermediate labeled-data tasks](#).
- 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.
- Stephen Roller, Emily Dinan, Naman Goyal, Da Ju, Mary Williamson, Yinhan Liu, Jing Xu, Myle Ott, Eric Michael Smith, Y-Lan Boureau, and Jason Weston. 2021. [Recipes for building an open-domain chatbot](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 300–325, Online. Association for Computational Linguistics.
- Wataru Sakata, Tomohide Shibata, Ribeka Tanaka, and Sadao Kurohashi. 2019. [Faq retrieval using query-question similarity and bert-based query-answer relevance](#).
- Kurt Shuster, Spencer Poff, Moya Chen, Douwe Kiela, and Jason Weston. 2021. [Retrieval augmentation reduces hallucination in conversation](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 3784–3803, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Noriko Tomuro and Steven Lytinen. 2004. Retrieval models and q and a learning with faq files. pages 183–202.
- Kexin Wang, Nandan Thakur, Nils Reimers, and Iryna Gurevych. 2021. [Gpl: Generative pseudo labeling for unsupervised domain adaptation of dense retrieval](#). *arXiv preprint arXiv:2112.07577*.
- Haode Zhang, Yuwei Zhang, Li-Ming Zhan, Jiaxin Chen, Guangyuan Shi, Xiao-Ming Wu, and Albert Y.S. Lam. 2021a. [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.
- Jianguo Zhang, Trung Bui, Seunghyun Yoon, Xiang Chen, Zhiwei Liu, Congying Xia, Quan Hung Tran, Walter Chang, and Philip Yu. 2021b. [Few-shot intent detection via contrastive pre-training and fine-tuning](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 1906–1912, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- 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.

Appendix

A Dataset details

The details regarding different datasets is given in Table 7. The dataset BANKING77 consists of 77 intents for 1 domain, CLINC150 has 150 intents across 10 domains and HWU64 has 64 intents across 21 domains. Curekart, Powerplay11 and SOFMattress have 28, 59 and 21 intents respectively. Powerplay11 is a challenging set with some intents having only 1 sample even in full dataset and median samples per intent being 3 in the subset variant. The total number of Query-Question (Q-Q) pairs generated for each dataset is also shown.

Datasets	Intents (Domains)	Min	Max	Median	Samples	q-Q Pairs
BANKING77	77 (1)	30	167	108	8.6K	-
BANKING77-5	77 (1)	5	5	5	385	73920
BANKING77-10	77 (1)	10	10	10	770	296065
CLINC150	150 (10)	100	100	100	15K	-
CLINC150-5	150 (10)	5	5	5	750	280875
CLINC150-10	150 (10)	10	10	10	1500	1124250
HWU64	64 (21)	33	159	156	8.9K	-
HWU64-5	64 (21)	5	5	5	320	51040
HWU64-10	64 (21)	10	10	10	640	204480
Curekart	28 (1)	3	95	14	600	179700
Curekart-Subset	28 (1)	2	72	8	413	85078
SOFmattress	21 (1)	9	34	12	328	53628
SOFmattress-Subset	21 (1)	3	18	7	180	16110
Powerplay	59 (1)	1	46	7	471	110685
Powerplay-Subset	59 (1)	1	24	3	261	33930

Table 7: Details of datasets used along with number of Q-Q pair generated. Min, Max and Median are on samples per intent