

# Gradual Fine-Tuning for Low-Resource Domain Adaptation

Haoran Xu<sup>1</sup>, Seth Ebner<sup>1</sup>, Mahsa Yarmohammadi<sup>1</sup>, Aaron Steven White<sup>2</sup>,  
Benjamin Van Durme<sup>1</sup>, Kenton Murray<sup>1</sup>

<sup>1</sup>Johns Hopkins University, <sup>2</sup>University of Rochester

{hxu64, seth, mahsa, vandurme, kenton}@jhu.edu,  
aaron.white@rochester.edu

## Abstract

Fine-tuning is known to improve NLP models by adapting an initial model trained on more plentiful but less domain-salient examples to data in a target domain. Such domain adaptation is typically done using one stage of fine-tuning. We demonstrate that gradually fine-tuning in a multi-stage process can yield substantial further gains and can be applied without modifying the model or learning objective.

## 1 Introduction

Domain adaptation is a technique for practical applications in which one wants to learn a model for a task in a particular domain with too few instances of in-domain data to directly learn a model. Common approaches for domain adaptation make use of fine-tuning (Dabre et al., 2019; Li and Specia, 2019; Imankulova et al., 2019), in which a model is pretrained on a large amount of out-of-domain but task-relevant data and then refined toward the target domain by subsequently training on in-domain data. This fine-tuning procedure is often performed in one stage: the pretrained model is trained on the in-domain data until convergence (Chu et al., 2017; Min et al., 2017a). We propose a *gradual* fine-tuning approach, in which a model is iteratively trained to convergence on data whose distribution progressively approaches that of the in-domain data. Intuitively, the model is eased toward the target domain rather than abruptly shifting to it.

Inspired by the general approach of curriculum learning of training a model on a trajectory from easier instances to more difficult instances (Bengio et al., 2009), we train a model on a sequence of datasets, each of which would be increasingly difficult to learn on its own due to its size. Each dataset in the sequence interpolates between the data in the previous iteration and the target domain data. We hypothesize that just as in curriculum learn-

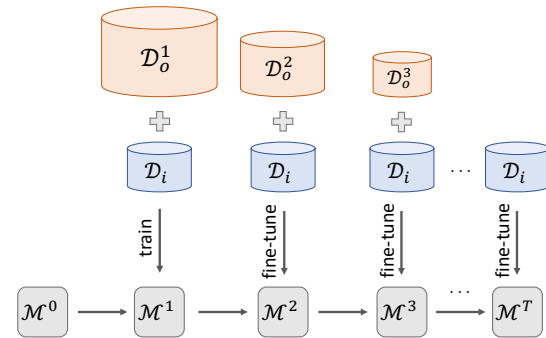


Figure 1: Stages of gradual fine-tuning: 1) Train the model,  $\mathcal{M}$ , on a mixture of in-domain data,  $\mathcal{D}_i$ , and out-of-domain data,  $\mathcal{D}_o$ ; 2) iteratively fine-tune on mixed domain data with decreasing amounts of out-of-domain data; 3) fine-tune on only in-domain data.

ing where first learning from easier instances helps models subsequently learn from more difficult instances, the interpolation process yields datasets with distributions that are increasingly useful for helping models learn during subsequent stages of the gradual fine-tuning procedure. We begin by training the model on data that contains a mix of out-of-domain and in-domain instances, then increase the concentration of in-domain data in each fine-tuning stage. In this way, at each stage of fine-tuning we increase the similarity between the current domain and the target domain, which enables the model to potentially better fit the distribution of the target domain. The approach is illustrated in Figure 1.

We conduct experiments on two NLP tasks to demonstrate the effectiveness of gradual fine-tuning. We first look at (1) dialogue state tracking with the MultiWOZ v2.0 dataset (Budzianowski et al., 2018), which is a collection of human-to-human conversation transcriptions in multiple domains. We focus on single dialogue domains and utilize out-of-domain data to improve accuracy of

slot classification in the target domain. We then consider (2) an event extraction task from the ACE 2005 dataset (LDC2006T06) for which we augment the Arabic target domain data with English data.

Gradual fine-tuning is also simple to implement. If fine-tuning is already supported by the code, then only new configuration files need to be created to specify the (amount of) data used at each iteration.<sup>1</sup> No adjustments to model or training code are needed. Just by modifying the training approach, one can obtain substantial improvements.<sup>2</sup>

## 2 Related Work

Howard and Ruder (2018) propose an effective inductive transfer learning method for language model fine-tuning and demonstrate improvements on text classification tasks. Gururangan et al. (2020) also show improvements on target task performance by fine-tuning pretrained language models on in-domain data and on the target task’s training data. In this work, we focus on adapting the entire model, not just the underlying language model encoder. Our approach is a form of transductive transfer (Pan and Yang, 2009), in which the pre-training and fine-tuning tasks are the same. In the transductive transfer setting, we hope to learn task-specific information by training on a large (potentially out-of-domain) dataset, and then subsequently adjust model parameters based on domain-specific information learned from in-domain data.

Transductive transfer has been effective for tasks such as question answering (Min et al., 2017b), machine translation (Sennrich et al., 2015), and open information extraction (Sarhan and Spruit, 2020). Wu et al. (2019) fine-tune toward a target domain for dialogue state tracking using Gradient Episodic Memory (Lopez-Paz and Ranzato, 2017) to avoid catastrophic forgetting (McCloskey and Cohen, 1989). Ahn et al. (2019) introduce an uncertainty-based regularization method to overcome catastrophic forgetting in the continual learning setting. There have also been successful approaches for cross-lingual information extraction and semantic role labeling using no target language data (Subburathinam et al., 2019), mixed source and translated target language data (Fei et al., 2020), and language-independent model transfer (Daza and

Frank, 2019; Fei et al., 2020).

Jiang and Zhai (2007) propose a method for upweighting the importance of target domain instances relative to source domain instances to improve domain adaptation. The iterative increase in concentration of target domain data in the mixed domain data used in gradual fine-tuning can be seen as analogous to giving target domain instances more weight. In contrast to all the aforementioned approaches, gradual fine-tuning requires no modification to existing models or learning objectives, so it can be applied to any system.

Domain adaptation can also be achieved using curriculum learning. Zhang et al. (2019) use a curriculum learning approach to adapt a general-domain machine translation model to a target domain while also using data whose domain is unknown. Inspired by curriculum learning (Bengio et al., 2009), which highlights the importance of the order of training instances, we propose a multi-stage fine-tuning strategy for domain adaptation. In this work, we order a sequence of fine-tuning datasets from least similar to the target domain to most similar.

## 3 Method

Domain adaptation via fine-tuning aims to improve performance on a target domain by using information learned from general domain data when training on data in the target domain. In other words, it is expected that training on general domain data would provide better initialization for subsequently training on target domain data than random initialization would provide (Erhan et al., 2010).

### 3.1 Mixed Domain Training

Given training data for the target domain,  $\mathcal{D}_i$ , we augment it with out-of-domain data  $\mathcal{D}_o$  mapped to the target task schema.<sup>3</sup> The in-domain and out-of-domain data are concatenated to form a mixed domain training set. A mixed domain model is obtained by training on the mixed dataset.

### 3.2 One-Stage Fine-tuning

The distribution of the mixed domain data is different from that of the target domain data. However, regardless of how diverse the mixed domain training data is, subsequently fine-tuning the model on

<sup>1</sup>One may also need to implement data subsampling.

<sup>2</sup>Supporting code: <https://github.com/felixxu/Gradual-Finetune>.

<sup>3</sup>Portions of the target task schema corresponding to fields not available in the out-of-domain data could be masked in the mapped data.

the target in-domain data is a direct way to encourage the model to converge to the distribution of the target domain. The model is expected to use task knowledge learned from the mixed domain data to yield improved performance over a model trained only on data from the target domain. Because the model is fine-tuned to convergence on the target domain data once, we refer to this procedure as one-stage fine-tuning.

---

### Algorithm 1 Gradual Fine-Tuning

---

**Require:** in-domain data  $\mathcal{D}_i$ , out-of-domain data  $\mathcal{D}_o^0$ , initial model  $\mathcal{M}^0$ , out-of-domain data schedule  $\mathcal{S}$

```

1: function GRADUAL-FT( $\mathcal{D}_i, \mathcal{D}_o^0, \mathcal{M}^0, \mathcal{S}$ )
2:    $t \leftarrow 0$ 
3:   for  $amount$  in  $\mathcal{S}$  do
4:      $t \leftarrow t + 1$ 
5:      $\mathcal{D}_o^t \leftarrow \text{SAMPLE}(\mathcal{D}_o^{t-1}, amount)$ 
6:      $\mathcal{D}_{train}^t \leftarrow \mathcal{D}_i \cup \mathcal{D}_o^t$ 
7:      $\mathcal{M}^t \leftarrow \text{TRAIN}(\mathcal{M}^{t-1}, \mathcal{D}_{train}^t)$ 
8:   end for
9:   return  $\mathcal{M}^t$ 
10: end function

```

---

### 3.3 Gradual Fine-tuning

Instead of adapting the model to the target domain by one-stage fine-tuning, we propose an iterative multi-stage approach that transitions from the initial mixed domain to the target domain. Each iteration incorporates less out-of-domain data than the preceding iteration as specified by a data schedule  $\mathcal{S}$ , bringing the data distribution closer to that of the target domain every training cycle. At each iteration, the model is trained to convergence. The number of iterations and the out-of-domain data schedule  $\mathcal{S}$  are hyperparameters that can be tuned for a particular task. Pseudocode for gradual fine-tuning is presented in Algorithm 1. In our implementation, the out-of-domain data is uniformly randomly sampled from the out-of-domain data used in the previous iteration.

## 4 Experiments

### 4.1 Dialogue State Tracking

Dialogue state tracking (DST) involves estimating at each dialogue turn the probability distribution over slot-values enumerated in an ontology. For example, we may be interested in the distribution over cuisines given the dialogue history and an utterance indicating the *restaurant* domain and food slot. We show that gradual fine-tuning can substantially improve slot accuracy—the accuracy of predicting

	Rest.	Hotel	Attract.	Taxi	Train
# Slots	7	10	3	4	6
# Turns	3011	3472	577	1667	1771
# Dialogues					
Train	523	513	127	326	282
Dev	50	56	11	57	30
Test	61	65	12	52	33

Table 1: Data statistics for five domains from MultiWOZ v2.0.

each slot separately—and joint accuracy—the percentage of turns in which all slots are predicted correctly—in a given dialogue domain.

**Dataset** We run experiments on the MultiWOZ v2.0 dataset (Budzianowski et al., 2018), which is a multi-domain conversational corpus with seven domains and 35 slots. Following Wu et al. (2019), we focus on five domains: *restaurant*, *hotel*, *attraction*, *taxi*, and *train*, which amounts to 2198 single-domain dialogues and 5459 multi-domain dialogues from the original dataset across all data splits. Statistics for single-domain dialogues of the five domains are presented in Table 1. Among the five domains, *restaurant* and *hotel* are adopted as the target domains for our experiments. We consider single-domain dialogues in the target domain as in-domain data and the rest of the (single-domain or multi-domain) dialogues excluding the target domain as out-of-domain data.

**Settings** The gradual fine-tuning data schedule of out-of-domain dialogues is  $\mathcal{S} = 4\text{K} \rightarrow 2\text{K} \rightarrow 0.5\text{K} \rightarrow 0$  (K=thousand), where the in-domain data is mixed with the out-of-domain data at each stage.<sup>4</sup> We use the Slot-Utterance Matching Belief Tracker (SUMBT) model by Lee et al. (2019), as well as their hyperparameters. The SUMBT model is composed of four parts: BERT encoders for encoding slots, values, and utterances, a slot-utterance matching network, a belief tracker, and a nonparametric discriminator. SUMBT achieves state-of-the-art performance on the MultiWOZ v2.0 dataset. More training details can be found in Appendix A.

**Baselines** Three baselines are considered in this experiment. The first one is a model trained only on in-domain data (no data augmentation). The second is a model trained with the one-stage fine-

<sup>4</sup>Experiments starting with 2K ( $2\text{K} \rightarrow 0.5\text{K} \rightarrow 0$ ) augmented data examples were also conducted to further explore the effect of data size (Appendix B).

tuning strategy ( $\mathcal{S} = 4\text{K} \rightarrow 0$ ). The last baseline is a model trained with the same settings as Lee et al. (2019), which has seen the full training set.<sup>5</sup>

**Result and Analysis** The main results are shown in Table 2. Compared with the model trained without data augmentation, gradual fine-tuning yields an absolute gain of 3.6% for slot accuracy and 15.11% for joint accuracy in the *restaurant* domain and 1.7% slot accuracy and 5.82% joint accuracy in the *hotel* domain. Moreover, gradual fine-tuning also considerably outperforms both the one-stage strategy as well as the mixed data training.

Figure 2 shows increasing slot accuracy at each stage of (gradual) fine-tuning, which supports our hypothesis that gradually fine-tuning the model consistently improves performance as the data distribution approaches that of the target domain.

	Restaurant Slot/Joint	Hotel Slot/Joint
No FT (single domain)	90.70/52.16	90.79/46.30
No FT* (all domains)	92.19/58.63	91.48/50.26
One-stage FT	93.47/61.15	91.43/46.30
Gradual FT	<b>94.30/67.27</b>	<b>92.49/52.12</b>

Table 2: Slot and joint accuracy in the restaurant and hotel domains under various training methods. \* indicates that the model uses the full training set. *No FT (single domain)* is the model trained only on in-domain data and *No FT (all domains)* is the training regime from Lee et al. (2019). FT=fine-tuning.

## 4.2 Event Extraction

We also employ gradual fine-tuning on an event extraction task to show the general applicability of the approach. Event extraction involves predicting event triggers, event arguments, and argument roles. We perform event extraction on the ACE 2005 corpus by considering Arabic as the target domain and English as the auxiliary domain.

**Settings** For data processing, model building, and performance evaluation, we use the DYGIE++ framework,<sup>6</sup> which achieved state-of-the-art results on ACE 2005 event extraction (Wadden et al., 2019). We replace the BERT encoder (Devlin et al., 2019) with XLM-R (Conneau et al., 2020) to train models on monolingual and mixed bilingual

<sup>5</sup>Note that use of the full training set means the model sees multi-domain dialogues which may include the target domain. The full training set contains approximately 4K more dialogues than what we use for gradual fine-tuning.

<sup>6</sup><https://github.com/dwadden/dygiepp>

datasets. Because no standard splits were found for the Arabic portion of the ACE 2005 dataset, the train/dev/test splits for Arabic are randomly selected. Table 3 shows statistics of the ACE 2005 events dataset for our experiments. The gradual fine-tuning data schedule of preprocessed English documents,  $\mathcal{S}$ , is  $1\text{K} \rightarrow 0.5\text{K} \rightarrow 0.2\text{K} \rightarrow 0$ .<sup>7</sup> We report four metrics for the evaluation: a trigger is correctly identified if its offsets find a match in the ground truth (TrigID), and it is correctly classified if their event types match (TrigC). An argument is correctly identified if its offsets and event type find a match in the ground truth (ArgID), and it is correctly classified if their event roles match (ArgC).

	English	Arabic
# Event types	33	30
# Role types	22	21
# Events/Arguments		
Train	4202/4859	1743/2506
Dev	450/605	117/174
Test	403/576	198/287

Table 3: Statistics for English and Arabic ACE 2005.


**Baselines** The first baseline in Table 4 is the model from Wadden et al. (2019) (with an XLM-R encoder) trained only on Arabic data, the second baseline is trained on mixed data (Arabic + 1K English data) without any fine-tuning, and the third baseline uses the one-stage fine-tuning strategy.

**Results** The results are given in Table 4. Mixed data training and one-stage fine-tuning achieve slight improvements over no data augmentation and even hurt performance on the TrigID and ArgID metrics. Gradual fine-tuning outperforms all baselines by a substantial margin on all four metrics, especially TrigC, which achieves an absolute gain of 5.84% compared to the state-of-the-art model trained only on the Arabic dataset.


	TrigID	TrigC	ArgID	ArgC
No FT (Ar)	64.77	57.03	47.76	42.83
No FT (mixed)	64.12	59.48	46.57	43.21
One-stage FT	63.61	59.88	46.79	43.44
Gradual FT	<b>66.29</b>	<b>62.87</b>	<b>48.11</b>	<b>44.21</b>

Table 4: Identification and classification F1 scores for triggers and arguments on Arabic ACE 2005. FT=fine-tuning, Ar=Arabic, mixed=Arabic + 1K English.

<sup>7</sup>1K, 0.5K, and 0.2K preprocessed documents correspond to 85%, 35%, and 5% of total events/arguments in the English training set.



(a) Restaurant



(b) Hotel

Figure 2: Slot accuracy for the *restaurant* and *hotel* domains in MultiWOZ v2.0. Gradual fine-tuning yields higher performance than fine-tuning in a one-stage process. Both fine-tuning approaches perform better than not doing any fine-tuning. Results at the end of training are the rightmost points (Table 2).

## 5 Conclusion

We have proposed a gradual fine-tuning technique that iteratively steers the distribution of augmented training data toward that of a target domain. Gradual fine-tuning can be straightforwardly applied to an existing codebase without changing the model architecture or learning objective. Through experiments on dialogue state tracking and event extraction tasks, we have demonstrated that gradual fine-tuning outperforms standard one-stage fine-tuning for domain adaptation.

## Acknowledgments

We thank the anonymous reviewers for their valuable comments. This work was supported in part by IARPA BETTER (#2019-19051600005) and DARPA KAIROS (FA8750-19-2-0034). The views and conclusions contained in this work are those of the authors and should not be interpreted as necessarily representing the official policies, either expressed or implied, or endorsements of DARPA, ODNI, IARPA, or the U.S. Government. The U.S. Government is authorized to reproduce and distribute reprints for governmental purposes notwithstanding any copyright annotation therein.

## References

- Hongjoon Ahn, Sungmin Cha, Donggyu Lee, and Taesup Moon. 2019. [Uncertainty-based continual learning with adaptive regularization](#). In *Advances in Neural Information Processing Systems*, volume 32. Curran Associates, Inc.
- Yoshua Bengio, Jérôme Louradour, Ronan Collobert, and Jason Weston. 2009. [Curriculum learning](#). In

*Proceedings of the 26th Annual International Conference on Machine Learning, ICML '09*, page 41–48, New York, NY, USA. Association for Computing Machinery.

- Paweł Budzianowski, Tsung-Hsien Wen, Bo-Hsiang Tseng, Iñigo Casanueva, Ultes Stefan, Ramadan Osman, and Milica Gašić. 2018. [MultiWOZ - A Large-Scale Multi-Domain Wizard-of-Oz Dataset for Task-Oriented Dialogue Modelling](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing (EMNLP)*.

- Chenhui Chu, Raj Dabre, and Sadao Kurohashi. 2017. [An empirical comparison of domain adaptation methods for neural machine translation](#). In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 385–391, Vancouver, Canada. Association for Computational Linguistics.

- Alexis Conneau, Kartikay Khandelwal, Naman Goyal, Vishrav Chaudhary, Guillaume Wenzek, Francisco Guzmán, Edouard Grave, Myle Ott, Luke Zettlemoyer, and Veselin Stoyanov. 2020. [Unsupervised cross-lingual representation learning at scale](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8440–8451, Online. Association for Computational Linguistics.

- Raj Dabre, Atsushi Fujita, and Chenhui Chu. 2019. [Exploiting multilingualism through multistage fine-tuning for low-resource neural machine translation](#). 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 1410–1416.

- Angel Daza and Anette Frank. 2019. [Translate and label! an encoder-decoder approach for cross-lingual semantic role labeling](#). 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 603–615, Hong Kong, China. 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*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Dumitru Erhan, Aaron Courville, Yoshua Bengio, and Pascal Vincent. 2010. Why does unsupervised pre-training help deep learning? In *Proceedings of the Thirteenth International Conference on Artificial Intelligence and Statistics*, pages 201–208.
- H. Fei, M. Zhang, F. Li, and D. Ji. 2020. Cross-lingual semantic role labeling with model transfer. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 28:2427–2437.
- Hao Fei, Meishan Zhang, and Donghong Ji. 2020. Cross-lingual semantic role labeling with high-quality translated training corpus. *arXiv preprint arXiv:2004.06295*.
- 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.
- Jeremy Howard and Sebastian Ruder. 2018. **Universal language model fine-tuning for text classification**. In *Proceedings of the 56th Annual Meeting of the Association for Computational Linguistics*, pages 328–339, Melbourne, Australia. Association for Computational Linguistics.
- Aizhan Imankulova, Raj Dabre, Atsushi Fujita, and Kenji Imamura. 2019. **Exploiting out-of-domain parallel data through multilingual transfer learning for low-resource neural machine translation**. In *Proceedings of Machine Translation Summit XVII Volume 1: Research Track*, pages 128–139, Dublin, Ireland. European Association for Machine Translation.
- Jing Jiang and ChengXiang Zhai. 2007. **Instance weighting for domain adaptation in NLP**. In *Proceedings of the 45th Annual Meeting of the Association of Computational Linguistics*, pages 264–271, Prague, Czech Republic. Association for Computational Linguistics.
- Hwaran Lee, Jinsik Lee, and Tae-Yoon Kim. 2019. **SUMBT: Slot-utterance matching for universal and scalable belief tracking**. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 5478–5483, Florence, Italy. Association for Computational Linguistics.
- Zhenhao Li and Lucia Specia. 2019. **Improving neural machine translation robustness via data augmentation: Beyond back-translation**. In *Proceedings of the 5th Workshop on Noisy User-generated Text (W-NUT 2019)*, pages 328–336, Hong Kong, China. Association for Computational Linguistics.
- David Lopez-Paz and Marc’Aurelio Ranzato. 2017. Gradient episodic memory for continual learning. In *Advances in neural information processing systems*, pages 6467–6476.
- Michael McCloskey and Neal J Cohen. 1989. Catastrophic interference in connectionist networks: The sequential learning problem. In *Psychology of learning and motivation*, volume 24, pages 109–165. Elsevier.
- Sewon Min, Minjoon Seo, and Hannaneh Hajishirzi. 2017a. **Question answering through transfer learning from large fine-grained supervision data**. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 2: Short Papers)*, pages 510–517, Vancouver, Canada. Association for Computational Linguistics.
- Sewon Min, Minjoon Seo, and Hannaneh Hajishirzi. 2017b. Question answering through transfer learning from large fine-grained supervision data. *arXiv preprint arXiv:1702.02171*.
- Sinno Jialin Pan and Qiang Yang. 2009. A survey on transfer learning. *IEEE Transactions on knowledge and data engineering*, 22(10):1345–1359.
- Injy Sarhan and Marco Spruit. 2020. Can we survive without labelled data in nlp? transfer learning for open information extraction. *Applied Sciences*, 10(17):5758.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2015. Improving neural machine translation models with monolingual data. *arXiv preprint arXiv:1511.06709*.
- Ananya Subburathinam, Di Lu, Heng Ji, Jonathan May, Shih-Fu Chang, Avirup Sil, and Clare Voss. 2019. **Cross-lingual structure transfer for relation and event extraction**. 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 313–325, Hong Kong, China. Association for Computational Linguistics.
- David Wadden, Ulme Wennberg, Yi Luan, and Hannaneh Hajishirzi. 2019. **Entity, relation, and event extraction with contextualized span representations**. 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 5784–5789, Hong Kong, China. Association for Computational Linguistics.

Chien-Sheng Wu, Andrea Madotto, Ehsan Hosseini-Asl, Caiming Xiong, Richard Socher, and Pascale Fung. 2019. [Transferable multi-domain state generator for task-oriented dialogue systems](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 808–819, Florence, Italy. Association for Computational Linguistics.

Xuan Zhang, Pamela Shapiro, Gaurav Kumar, Paul McNamee, Marine Carpuat, and Kevin Duh. 2019. [Curriculum learning for domain adaptation in neural machine translation](#). 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 1903–1915, Minneapolis, Minnesota. Association for Computational Linguistics.

## A Learning Rate Schedule

### A.1 Potential Model Collapse

If we use the same constant learning rate for each fine-tuning stage, the model may collapse. Accuracy on the train and dev sets may drop to 0 or near-zero at the beginning of a stage of gradual fine-tuning. This phenomenon is often caused by using an overly large learning rate. This suggests that one must carefully schedule the learning rates for each stage of gradual fine-tuning.

### A.2 Learning Rates for MultiWOZ v2.0

In the MultiWOZ v2.0 experiments, we use  $4e-5$  for the learning rate in the last stage of fine-tuning and  $1e-4$  (default settings from Lee et al. (2019)) for the other stages.

### A.3 Learning Rates for ACE 2005

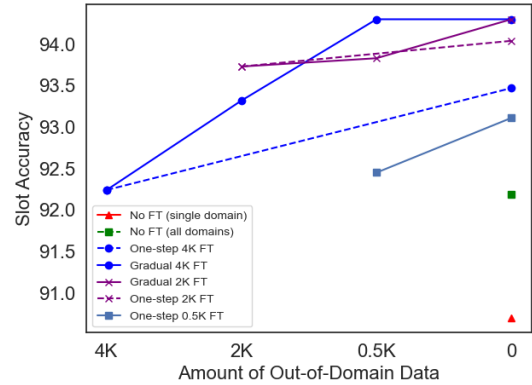
The network is split into 2 parameter groups: the parameters of the XLM-R encoder and all other parameters. We set the base learning rate of XLM-R to  $5e-5$  and to  $1e-3$  for the other parameters (default settings from Wadden et al. (2019)). We do not change the base learning rate for the first two stages of fine-tuning. At the third stage, we reduce the learning rate of XLM-R to  $1e-5$  and of other parameters to  $4e-4$ . We further reduce the learning rate of XLM-R to  $8e-6$  and the all other learning rates to  $2e-4$  for the final stage of gradual fine-tuning.

## B Additional MultiWOZ v2.0 Results


Here we present additional results of gradual fine-tuning on MultiWOZ v2.0 to augment Table 2. Table 5 shows the results of one-stage fine-tuning and gradual fine-tuning using the data schedule  $\mathcal{S} = 4K \rightarrow 2K \rightarrow 0.5K \rightarrow 0$  as well as suffixes of  $\mathcal{S}$ . Figure 3 shows the trend of slot accuracy for all training strategies at each stage of fine-tuning.

	Restaurant Slot/Joint	Hotel Slot/Joint
No FT (single domain)	90.70/52.16	90.79/46.30
No FT* (all domains)	92.19/58.63	91.48/50.26
One-stage 0.5k FT	93.11/64.75	91.61/46.30
One-stage 2k FT	94.04/62.23	91.24/47.35
One-stage 4k FT	93.47/61.15	91.43/46.30
Gradual 2k FT	94.30/66.91	91.35/46.30
Gradual 4k FT	<b>94.30/67.27</b>	<b>92.49/52.12</b>

Table 5: Slot and joint accuracy in the *restaurant* and *hotel* domains under various training methods. \* indicates that the model uses the full training set. FT=fine-tuning.



(a) Restaurant



(b) Hotel

Figure 3: Slot accuracy for the *restaurant* and *hotel* domains in MultiWOZ v2.0.