

ILLUMINER: Instruction-tuned Large Language Models as Few-shot Intent Classifier and Slot Filler

Paramita Mirza*, Viju Sudhi†, Soumya Ranjan Sahoo†,
Sinchana Ramakanth Bhat†

*Fraunhofer IIS, †Fraunhofer IAIS

paramita.paramita@iis.fraunhofer.de,

{viju.sudhi, soumya.ranjan.sahoo, sinchana.ramakanth.bhat}@iais.fraunhofer.de

Abstract

State-of-the-art intent classification (IC) and slot filling (SF) methods often rely on data-intensive deep learning models, limiting their practicality for industry applications. Large language models on the other hand, particularly instruction-tuned models (Instruct-LLMs), exhibit remarkable zero-shot performance across various natural language tasks. This study evaluates Instruct-LLMs on popular benchmark datasets for IC and SF, emphasizing their capacity to learn from fewer examples. We introduce ILLUMINER, an approach framing IC and SF as language generation tasks for Instruct-LLMs, with a more efficient SF-prompting method compared to prior work. A comprehensive comparison with multiple baselines shows that our approach, using the FLAN-T5 11B model, outperforms the state-of-the-art joint IC+SF method and in-context learning with GPT3.5 (175B), particularly in slot filling by 11.1–32.2 percentage points. Additionally, our in-depth ablation study demonstrates that parameter-efficient fine-tuning requires less than 6% of training data to yield comparable performance with traditional full-weight fine-tuning.

Keywords: intent classification, slot filling, instruction-tuned models, parameter-efficient fine-tuning

1. Introduction

Intent classification (IC) and slot filling (SF) are foundational tasks in natural language understanding (NLU) within task-oriented dialogue (TOD) systems, which enable users to interact in natural language, facilitating various actions such as reserving a restaurant or seeking customer support. For instance, given a user utterance “Find me a restaurant serving Italian food in Torino”, IC discerns the user’s intent as *find restaurant*, while SF aims for extracting slot type–value pairs $\{(cuisine, 'Italian'), (city, 'Torino')\}$ from the utterance. This information is crucial for generating appropriate system responses. Furthermore, efficiently and reliably solving these tasks with low latency is vital for the widespread deployment of TOD systems. Although deep learning models have excelled in supervised learning approaches (Gupta et al., 2019; Chen et al., 2019; Han et al., 2022), their reliance on large-scale annotated data constrains their practical use in real-world industrial scenarios.

Large language models (LLMs), especially those fine-tuned with instructions (Instruct-LLMs), have been touted as effective zero-shot learners (Wei et al., 2022). Instruction tuning empowers these models to interpret and execute user instructions effectively, thereby controlling their behavior (Zhang et al., 2023). Unlike supervised fine-tuning, which relies on input examples and their corresponding outputs, instruction tuning augments input–output examples with *instructions* as high-level task descriptions (depicted in Figure 1). This allows instruction-tuned models to generalize more read-

ily to new tasks or domains. When combined with *in-context learning* (Brown et al., 2020), where the model is exposed to input–output examples within the *prompt*, LLM-prompting methods offer substantial benefits over traditional supervised approaches in terms of reduced labeled data requirements. In-context learning (ICL), or *few-shot learning*, offers language models a chance to learn from examples, but models’ context size often limits the number of examples. Processing k training examples for k -shot ICL also increases inference time k times as the prompt size grows (Liu et al., 2022b). While fine-tuning LLMs with more examples from downstream datasets yields substantial performance gains compared to using them out-of-the-box (Su et al., 2022; Xie et al., 2022), full-weight fine-tuning on consumer hardware is impractical and risks *catastrophic forgetting* (Goodfellow et al., 2015), particularly when the downstream dataset is small and lacks diversity. *Parameter-efficient fine-tuning* (PEFT, e.g. Hu et al. 2022a; Liu et al. 2022b) alleviates these issues by allowing fine-tuning of a small number of additional parameters while freezing most LLM parameters, significantly reducing computational and storage costs while retaining the LLMs’ prior, generalized knowledge.

Approach and Contributions. In this work, we introduce our approach **ILLUMINER**¹, Instruction-tuned Large Language Models as **IN**tent Classifier and Slot **FI**ller. We formulate IC and SF as language generation tasks, as exemplified in Figure 1.

¹<https://github.com/OpenGPTX/illuminer>

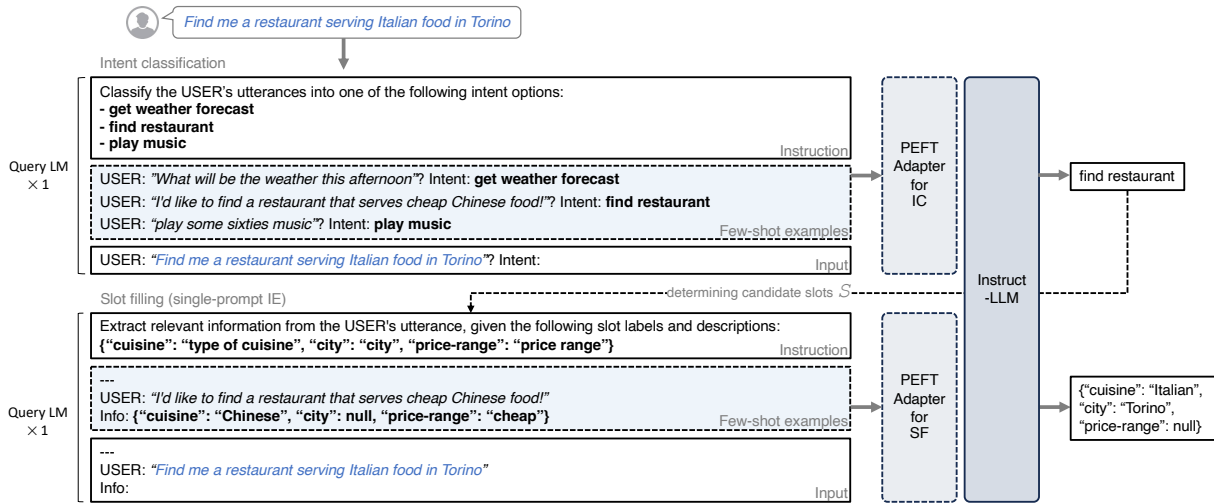


Figure 1: An example of our prompting methods for intent classification and slot filling, for a given user utterance “Find me a restaurant serving Italian food in Torino”. Compared to prior work (Fig. 2), we only need a single inference for slot filling.

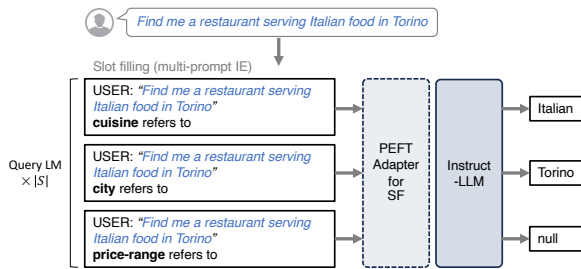


Figure 2: Multi-prompt IE for slot filling (Hou et al., 2022) requiring $|S|$ inferences for $|S|$ slot types.

For intent classification, we list possible intent labels to choose from in the instruction, and expect the Instruct-LLM to generate the appropriate label reflecting the intent of the input utterance. As opposed to prior work on slot filling with multiple prompts (Hou et al., 2022) illustrated in Figure 2, we adopt a *single-prompt Information Extraction (IE)* approach, requiring a single query per utterance for the Instruct-LLM to generate slot type–value pairs. We explore the performance of Instruct-LLMs further fine-tuned with task-specific instructions and domain-specific examples using PEFT approaches like *Low-Rank Adaptation (LoRA)* (Hu et al. 2022a) and *Infused Adapter by Inhibiting and Amplifying Inner Activations (IA)*³ by Liu et al. (2022b).

The salient contributions of our work are:

- Exploratory analysis of prompt engineering for IC and SF, and a much more efficient SF-prompting method compared to existing techniques (e.g., Hou et al. 2022; Li et al. 2023b).
- Comprehensive comparative analysis of several Instruct-LLMs on popular benchmark datasets for

IC and SF, including SNIPS (Coucke et al., 2018), MASSIVE (FitzGerald et al., 2022) and MultiWoz (Budzianowski et al., 2018), in various settings: zero-shot learning, few-shot learning, and PEFT. We demonstrate that ILLUMINER (with PEFT) outperforms state-of-the-art baselines, particularly in SF, given less than 6% of training data.

- Extensive ablation study examining the impact of Instruct-LLMs, different PEFT techniques, model size, number of examples for fine-tuning, and label exposure in instructions, as well as generalization across datasets.

2. Related Work

PLMs for IC and SF. The rise of pre-trained language models (PLMs) like BERT (Devlin et al., 2019) and RoBERTa (Liu et al., 2019) has spurred extensive research in utilizing contextual embeddings for sequence classification and labeling, notably in the joint task of IC and SF (Gupta et al., 2019; Chen et al., 2019; Han et al., 2022), as has been well-documented by Weld et al. (2023). A joint IC+SF model offers the advantage of training/fine-tuning a single model while capitalizing on label correlations between intents and slots. However, it demands a large annotated corpus (Weld et al., 2023), rendering its application impractical in real-world scenarios.

Addressing the few-shot scenarios of IC and SF, existing work explores PLMs from three main perspectives: (1) *task-adaptive fine-tuning* (Zhang et al., 2021; Yu et al., 2021; Ma et al., 2021; Hou et al., 2021), (2) *data augmentation* (Rosenbaum et al., 2022; Lin et al., 2023), and (3) *prompt-based learning* (Hou et al., 2022; Parikh et al., 2023). Our work aligns with prompt-based learning, a crucial

consideration in low-resource scenarios where fine-tuning large PLMs (LLMs) is not feasible (Radford et al., 2019; Schick and Schütze, 2021). Limited studies focus on prompt-based learning for IC or SF. Parikh et al. (2023) explore different zero- and few-shot methods for IC, including LLM prompting and parameter-efficient fine-tuning. Hou et al. (2022) introduce a multi-prompt method for SF (Fig. 2), accelerating inference compared to the classic prompting that requires inference for every n-gram word span (Cui et al., 2021). Yet, a comprehensive evaluation of prompt-based learning for both IC and SF jointly is lacking.

While previous studies (Wang et al., 2022a; Hu et al., 2022b; Gupta et al., 2022; Hudeček and Dušek, 2023) explore prompt-based learning for the dialog state tracking (DST) task, they mostly focus on the dialog-level performance, making it difficult to analyze the performance on single dialog turns and on the specific sub-tasks of DST: IC and SF. Bridging this gap between the two lines of work (IC+SF vs DST) and demonstrating our approach’s generality, we also evaluate ILLUMINER on MultiWoz (Budzianowski et al., 2018), a prominent benchmark dataset for DST.

Parameter-efficient Fine-tuning (PEFT). Fine tuning large language models is often a compute-intensive task, demanding immoderate time, cost and monitoring resources. However, recent advances in PEFT address these challenges by learning considerably fewer LLM parameters while achieving comparable performance to models with fully fine-tuned weights. Notable PEFT techniques include adapter tuning (Houlsby et al., 2019), prefix tuning (Li and Liang, 2021), prompt tuning (Lester et al., 2021), LoRA (Hu et al., 2022a) and (IA)³ (Liu et al., 2022a).

In the realm of task-oriented dialogue systems, PEFT has been explored for IC, SF and response generation. Hung et al. (2022) train adapters for individual domains, demonstrating their composition for multi-domain specialization. Wang et al. (2022b) use adapter tuning with a copy network to prevent catastrophic forgetting and ensures entity consistency in dialogue flow. Fuisz et al. (2022) employ lightweight adapters on QA-tuned PLMs for SF treated as a Question Answering task. Li et al. (2023b) explore prefix tuning for cross-domain SF, while Chang et al. (2023) demonstrate the use of prompt tuning for IC and SF with speech models. Kwon et al. (2023) show that multilingual mT0 models fine-tuned with LoRA outperform baselines on IC and SF sub-tasks for low-resource languages. Additionally, Parikh et al. (2023) demonstrate how FLAN-T5, fine-tuned with (IA)³ adapters, outperforms larger language models like GPT-3 in IC.

Existing work examined different PEFT techniques

individually on benchmark datasets. To the best of our knowledge, our work is the first to compare and contrast different PEFT techniques for fine-tuning LLMs for IC and SF, including an exploration of cross-dataset generalization offered by each technique.

Instruction Tuning (IT). Firstly introduced by Wei et al. (2022), IT explores language models’ cross-task generalization through supervised fine-tuning with task-specific instructions and desired output (Zhang et al., 2023). Aligning the next-word prediction with user instructions enhances control and predictability, gaining traction with models like InstructGPT (Ouyang et al., 2022) and FLAN-T5 (Wei et al., 2022), which oftentimes outperform their respective base models. Constructing instruction datasets often involves using templates to transform text–label pairs into instruction–output pairs (Muennighoff et al., 2022; Longpre et al., 2023). Our approach adopts this method for constructing IC and SF datasets for fine-tuning task-specific adapters with PEFT.

Despite prior work on IT, its impact on NLU tasks like intent classification and slot filling has been under-explored. Our work aligns with LINGUIST (Rosenbaum et al., 2022), which focuses on generating annotated data for IC and SF labels through instruction tuning. The generated data, however, is used to fine-tune a BERT-based model for joint IC+SF (Chen et al., 2019), while we fine-tune Instruct-LLMs for IC- and SF-prompting.

3. Methodology

Problem Statement. We consider two NLU tasks where a single user utterance \mathbf{x} , with tokens x_1, x_2, \dots, x_n , yields an output structure \mathbf{y} . For example, given \mathbf{x} = “Find me an Italian restaurant with a parking lot”, \mathbf{y} in intent classification (IC) is an intent label l (e.g., *find restaurant*) representing the user’s intent in \mathbf{x} . For slot filling (SF), \mathbf{y} comprises slot type–value pairs $\{(t_i, v_i)\}_{i=1}^m$, with t_i as the slot type (e.g., *cuisine*) and v_i as the corresponding slot value (e.g., *Italian*) extracted from \mathbf{x} . In contrast to traditional slot filling where v_i is always a span of \mathbf{x} , we also consider slots where v_i can be inferred from \mathbf{x} , e.g., (*parking available*, *yes*) from “...with a parking lot”. This scenario is present in dialog state tracking (DST) benchmark datasets, reflecting a more realistic downstream application.

ILLUMINER, our approach for IC and SF using instruction-tuned LLMs, is illustrated in Figure 1. The input utterance \mathbf{x} is transformed into a task-specific prompt for either IC or SF. This prompt is then provided to an Instruct-LLM, possibly enhanced by a task-specific PEFT adapter. Conditioned on the prompt, the Instruct-LLM generates \mathbf{y} specific to the task.

P_1	instruction	Given the possible intents: $\{L\}$
	input	What is the user's intent in 'x'? Intent:
P_2	instruction	Given the following options: $\{L\}$
	input	What did the user want when the user said, 'x'? Answer:
P_3	instruction	Classify the USER's utterances into one of the following intent options: $\{L\}$
	input	USER: 'x' Intent:
P_4	instruction	Given a USER's utterance, choose one of the following intents: $\{L\}$
	input	USER: 'x' Intent:

Table 1: Prompt template variations for IC.

A prompt typically consists of an *instruction* distinguishing IC from SF prompts, and an *input* containing \mathbf{x} . For the in-context learning approach, the prompt also contains *few-shot examples* between the *instruction* and the *input*, in which each example utterance \mathbf{x}' follows the same template as \mathbf{x} , but is accompanied by the expected \mathbf{y}' . Note that providing few-shot examples to prompt an Instruct-LLM already enhanced with a PEFT adapter does not necessarily enhance performance and leads to longer inference time due to extended prompts; thus, they are never utilized in conjunction.

Prompt Engineering for IC. We include the list of possible intent labels L in the *instruction*, derived from the ground-truth intents in the evaluation set of considered datasets. Instead of the original intent labels as annotated in the dataset, we employ handcrafted intent descriptions as labels, e.g., ‘turn light on’ (*iot_hue_lighton*), ‘express liking music’ (*music_likeness*), as they enhance label semantics and improve Instruct-LLMs’ comprehension. We explore four prompt template variations for IC (Table 1), with L listed in a single line per label.

Prompt Engineering for SF. In the *instruction* for SF, we expose the list of candidate slots S in the form of $\{t_i: d_i\}_{i=1}^m$ where t_i is a slot type (e.g., *cuisine*) and d_i is its corresponding description (e.g., ‘type of cuisine’). Candidate slots S are those relevant for the user’s intent in a given utterance, e.g., *cuisine* and *price-range* for the *find restaurant* intent. We derived relevant slot types based on intent–slot type co-occurrences (at least once) in the training data. When constructing \mathbf{y}' for the *few-shot examples*, we insert *null* for relevant slots not present in the ground truth slots. For example, with $\mathbf{x}' = \text{“I’d like to find a restaurant that serves Chinese food!”}$, $\mathbf{y}' = \{(cuisine, \text{‘Chinese’}), (price-range, null)\}$.

In datasets like SNIPS and MASSIVE, annotations include general slot types like *time* and *city*. We designate general slots S_G as slot types co-occurring with more than three intent labels. We incorporate S_G alongside relevant slots S as part of the *instruction*. Regarding prompt templates, we explore one variation illustrated in Figure 1, following several iterations of prompt design in a preliminary study.

Dataset	# Intents	# Slots	Avg. prompt length			
			zero-shot		few-shot	
			IC	SF	IC	SF
SNIPS	7	45	75.4	115.2	294.4	493.8
MASSIVE	60	55	336.7	160.5	603.6	551.9
MultiWoz	11	24	83.9	90.8	450.9	303.5

Table 2: Datasets for IC and SF experiments.

4. Experimental Setup

Dataset. We consider (i) *SNIPS* (Coucke et al., 2018), (ii) *MASSIVE* (FitzGerald et al., 2022) (English split) and (iii) *MultiWoz 2.2* (Budzianowski et al., 2018) as our benchmark datasets (see Table 2) since they encompass both IC and SF objectives and are widely used in the community. For MultiWoz, we consider only the first turn of each conversation in the test set for evaluation, as the first turn typically conveys a clear intent and precise slots in a single utterance, while subsequent turns may necessitate dialogue history for context.

Evaluation Metric. Following the baselines, we evaluate the performance of our proposed approach using the standard automatic evaluation metrics of *accuracy* for IC and *micro F1-score* for SF. We define *hallucinations* for IC and SF as the ratio of false positives that cannot be found in candidate intent/slot labels and user utterances.

Models. We explore various Instruct-LLMs:

- *Falcon-7B-Instruct* (tiiuae/falcon-7b-instruct), Falcon-7B (Almazrouei et al., 2023) fine-tuned on a mixture of chat/instruct datasets (Penedo et al., 2023; Xu et al., 2023b).
- *BLOOMZ* (bigscience/bloomz-7b1), fine-tuned BLOOM (Workshop et al., 2022) on xP3 (Muenighoff et al., 2022), a collection of human-instruction datasets in 46 languages.
- *FLAN-T5* (google/flan-t5-xxl), fine-tuned T5 (11B) on the Flan Collection (Longpre et al., 2023), <instruction, output> pairs constructed from 62 datasets of 12 NLP tasks.
- *Vicuna* (lmsys/vicuna-13b-v1.5), from fine-tuning LLaMA 2 (13B, Touvron et al. 2023a) on 70K user-shared conversations collected from a website.
- *WizardLM* (WizardLM/WizardLM-13B-V1.1), fine-tuned LLaMA (13B, Touvron et al. 2023b) on the Evol-Instruct dataset (Xu et al., 2023a).

We chose medium-sized LLMs (7B–13B) to compare various Instruct-LLMs of similar size but with different architectures and fine-tuned on distinct datasets. To restrict the generation, we set 10 as the maximum new tokens for IC and 100 for SF.

Zero-shot vs Few-shot. In the few-shot setting, prompts contain k examples of user utterances and desired outputs (i.e., intent labels or slot type–value pairs), in contrast to the zero-shot setting.

For intent classification, we randomly select one example per intent label from a small training set (k examples per label, where $k = 10$), forming the few-shot set F . Due to the limited context size of LLMs, we randomly sample 10 examples from F when necessary. For slot filling, we randomly sample utterances from the training data until we fulfill the requirement of one example per slot type, yielding the few-shot set F , which we also constrain to a size of 10. Table 2 details the average prompt length for IC and SF in zero- vs few-shot settings.

Parameter-efficient Fine-Tuning. We explore various PEFT techniques including prefix-tuning, prompt-tuning, LoRA and (IA)³, implemented by Hugging Face². We train the PEFT adapters separately for IC and SF on each dataset using a small training set (k examples per label, where $k = 10$). Adapters for IC are fine-tuned by varying the prompt templates as listed in Table 1. In both prefix and prompt tuning, we learn 20 virtual tokens with a learning rate of $1e-2$. In the case of LoRA, we set the hyper-parameters as $r = 16$, $\text{lora_alpha} = 32$ and $\text{lora_dropout} = 0.1$. We optimize the learning rate for Lora and (IA)³ by selecting from $\{5e-4, 1e-3, 5e-3\}$, and the number of epochs from $\{5, 10, 20\}$. We employ *AdamW* as the optimizer with its default hyper-parameters.

Baselines. We consider the following baselines for IC and SF tasks:

- *JointBERT*, BERT-based models (110M) for joint IC+SF (Chen et al., 2019).³ Using default hyperparameters (batch size of 32, learning rate of $5e-5$), we train models for 20 epochs in experiments with full training data. In experiments with a small training set ($k = 10$ per label), we reduced the batch size to 8 and train models for 50 epochs.
- OpenAI *GPT3.5* (text-davinci-003), 175B GPT3 (Brown et al., 2020) that has been trained on a larger dataset, enhancing its capability on understanding natural language instructions.
- *LINGUIST* (Rosenbaum et al., 2022), which leverages instruction-tuned LLMs (AlexaTM-5B) to generate annotated data for IC and SF, given few-shot examples as seed set. The generated data is used to fine-tune a BERT-style model for joint IC+ST (Chen et al., 2019).

5. Results and Analysis

Intent Classification. Table 3 summarizes the performance of various Instruct-LLMs across different settings for all datasets considered. FLAN-T5 (flan-t5-xxl) consistently outperforms other models,

including slightly larger ones like Vicuna (vicuna-13b-v1.5) and WizardLM (WizardLM-13B-V1.1). We suggest that for intent classification, encoder-decoder models, such as FLAN-T5, excel in capturing utterance meaning, leading to superior sequence classification performance. LoRA fine-tuning outperforms few-shot learning in most cases, especially for 7B models and on MASSIVE where the set of intent labels is significantly larger. Fine-tuning 7B models yields comparable results to larger models on SNIPS and MultiWoz, but the gap widens on the challenging MASSIVE. Examining standard deviation across different prompt templates (Table 1), FLAN-T5 emerges as the most robust model, with a standard deviation ≤ 0.01 . Notably, fine-tuning with LoRA also helps in reducing performance variance when using different prompts.

Slot Filling. We evaluate Instruct-LLMs for slot filling, taking into account ground truth intent labels in the prompt construction and few-shot example generation. As reported in Table 4, FLAN-T5, Vicuna and WizardLM exhibit competitive performance, with no clear winner. However, in the few-shot setting, Vicuna outperforms FLAN-T5 on MultiWoz whereas WizardLM on SNIPS and MASSIVE datasets, suggesting that causal decoder models excel when the generation capability is essential for producing structured outputs like slot type-value pairs. Fine-tuning with LoRA significantly improves performance by reducing false positives, i.e., the models learn how and when to fill slots with *null* values. Furthermore, the few-shot setting and LoRA mitigate hallucinations (indicated by numbers inside parentheses) considerably, as they help controlling the models' behavior to only fill the slots with relevant information found in the input utterances. In a comparison with prior work (Hou et al., 2022), where FLAN-T5, Vicuna and WizardLM are prompted using the multi-prompt IE strategy, we find that the best results from this technique (Table 4, underlined) are inferior to the best results from our prompting approach (Table 4, bold) for slot filling across all settings and datasets.

Comparison with Baselines. Here we consider the task of joint IC and SF, where the predicted intents are used for building the prompt for SF (determining candidate slots S), differing from previous SF experiments that used ground-truth intent labels. To address potential error propagation from IC to SF, we selected `flan-t5-xxlLoRA` to instantiate ILLUMINER, given its superior performance in IC (Table 3), coupled with the IC prompt template P_1 . We present its performance against considered baselines (§ 4) in Table 5.

Small LMs (e.g., BERT, 110M) outperform larger models when fine-tuned on full training data, as indicated in bold in Table 5. However, with a

²<https://github.com/huggingface/peft>

³<https://github.com/monologg/JointBERT>

Instruct-LLM	Size	SNIPS			MASSIVE			MultiWoz		
		zero-shot	few-shot	LoRA	zero-shot	few-shot	LoRA	zero-shot	few-shot	LoRA
falcon-7b-instruct	7B	.301 ±.15	.570 ±.03	.779 ±.06	.103 ±.05	.360 ±.03	.546 ±.00	.558 ±.38	.748 ±.01	.941 ±.02
bloomz-7b1	7B	.795 ±.08	.686 ±.10	.930 ±.01	.265 ±.06	.435 ±.02	.657 ±.01	.899 ±.02	.894 ±.06	.941 ±.01
flan-t5-xxl	11B	.937 ±.01	.940 ±.00	.962 ±.00	.726 ±.01	.741 ±.01	.825 ±.01	.973 ±.00	.982 ±.00	.979 ±.00
vicuna-13b-v1.5	13B	.574 ±.30	.920 ±.01	.950 ±.01	.333 ±.23	.688 ±.01	.759 ±.01	.425 ±.17	.977 ±.00	.972 ±.01
WizardLM-13B-V1.1	13B	.720 ±.25	.674 ±.20	.921 ±.01	.355 ±.11	.678 ±.01	.731 ±.01	.962 ±.02	.933 ±.02	.956 ±.01

Table 3: Intent accuracy in zero-shot, few-shot and LoRA settings, across different datasets and Instruct-LLMs. Numbers following \pm indicate standard deviation across different prompts.

Instruct-LLM	Size	SNIPS			MASSIVE			MultiWoz		
		zero-shot	few-shot	LoRA	zero-shot	few-shot	LoRA	zero-shot	few-shot	LoRA
ILLUMINER: Single-prompt IE										
falcon-7b-instruct	7B	.136 (.70)	.543 (.18)	.835 (.01)	.042 (.87)	.421 (.29)	.585 (.02)	.319 (.66)	.640 (.24)	.928 (.02)
bloomz-7b1	7B	.177 (.66)	.541 (.19)	.876 (.00)	.043 (.81)	.349 (.30)	.640 (.01)	.278 (.70)	.527 (.20)	.943 (.01)
flan-t5-xxl	11B	.310 (.35)	.647 (.14)	.909 (.01)	.125 (.37)	.473 (.21)	.735 (.00)	.462 (.46)	.753 (.18)	.945 (.02)
vicuna-13b-v1.5	13B	.222 (.43)	.554 (.08)	.908 (.01)	.103 (.59)	.369 (.14)	.724 (.03)	.500 (.36)	.859 (.10)	.957 (.01)
WizardLM-13B-V1.1	13B	.298 (.53)	.685 (.12)	.899 (.00)	.116 (.67)	.474 (.19)	.710 (.01)	.428 (.53)	.830 (.10)	.951 (.02)
Multi-prompt IE (Hou et al., 2022)										
flan-t5-xxl	11B	.221 (.54)	.380 (.22)	.904 (.01)	.058 (.74)	.195 (.45)	.658 (.00)	.360 (.47)	.547 (.31)	.933 (.02)
vicuna-13b-v1.5	13B	.111 (.73)	<u>.569 (.06)</u>	.755 (.01)	.021 (.90)	<u>.252 (.29)</u>	.597 (.04)	.146 (.89)	<u>.798 (.09)</u>	.889 (.05)
WizardLM-13B-V1.1	13B	.127 (.69)	.531 (.13)	.703 (.01)	.020 (.93)	.202 (.42)	.509 (.01)	.255 (.76)	.717 (.17)	.855 (.08)

Table 4: Slot filling F1 in zero-shot, few-shot and LoRA settings, across different datasets and Instruct-LLMs. Numbers inside parentheses indicate the ratio of wrong predictions caused by hallucinations.

Model	SNIPS		MASSIVE		MultiWoz	
	IC	SF	IC	SF	IC	SF
$k = 10$ per label						
ILLUMINER (flan-t5-xxl _{LoRA})	<u>.961</u>	<u>.899</u>	<u>.833</u>	<u>.720</u>	.978	<u>.946</u>
ILLUMINER (flan-t5-xxl _{few-shot})	.918	.600	.718	.440	.970	.746
JointBERT (Chen et al., 2019)	.907	.608	.718	.609	.958	.747
GPT3.5 _{zero-shot}	.913	.487	.716	.372	<u>.979</u>	.696
GPT3.5 _{few-shot}	.931	.633	.757	.398	.973	.831
Rosenbaum et al. (2022) †	.920	.823	-	-	-	-
Full training set						
ILLUMINER (flan-t5-xxl _{LoRA})	.967	.948	.871	.797	.989	.962
JointBERT (Chen et al., 2019)	.983	.965	.885	.797	.990	.834

Table 5: Comparison with baselines in terms of intent accuracy (IC) and slot filling F1 (SF). † denotes that numbers were taken directly from the paper.

small training set ($k = 10$ per label, 0.5%–5.2% of the full training set), ILLUMINER (flan-t5-xxl_{LoRA}) offers clear advantages over all baselines, delivering the best performance (underlined in Table 5), particularly on challenging datasets like MASSIVE (with 60 intent labels) and intricate tasks like slot filling. Medium-sized fine-tuned Instruct-LLMs (e.g., flan-t5-xxl_{LoRA}, 11B) even outperform few-shot learning with much larger models (e.g., GPT3.5_{few-shot}, 175B), highlighting the significance of exposing LLMs to more comprehensive examples, even if limited in size, achievable with parameter-efficient fine-tuning (PEFT). Provided with the same few-shot examples as GPT3.5_{few-shot}, ILLUMINER (flan-t5-xxl_{few-shot}) exhibits lower performance in most cases, although it remains comparable especially for intent classification.

LINGUIST (Rosenbaum et al., 2022) performs




Figure 3: Instruct-LLMs vs their corresponding base models (non-instruct).

similarly on SNIPS with its data augmentation method. Yet, approaches relying on *sequence labelling/tagging* for SF (e.g., JointBERT, LINGUIST) have limited capabilities on MultiWoz, where 13.8% of slot values are *inferred* from input utterances, capping recall at 0.862. This underscores the superiority of our SF-prompting method with ILLUMINER.

6. Ablation Studies

To better study and analyze the effectiveness of task-adapted instruction tuning with PEFT, we conduct the following series of ablation experiments:

Instruct- vs Non-instruct LLMs. For this study, we examine FLAN-T5-large (780M), BLOOMZ (7B) and Falcon-Instruct (7B) as Instruct-LLMs, with T5-large, BLOOM and Falcon as their non-instruct counterparts. We fine-tune and evaluate LoRA adapters for these models, and present the results in Figure 3. Our observations indicate that FLAN-T5 consistently exhibits superior performance over T5 in both tasks, showcasing its adept learning of




Figure 4: Performance of FLAN-T5_{LoRA} with various FLAN-T5 size. Solid-colored bars indicate adapters' training time for IC and striped bars for SF.




Figure 5: Performance of FLAN-T5-xxl_{LoRA} with various number of examples (k) per label. Solid-colored bars indicate % of training data for IC and striped bars for SF.

task-specific instructions. BLOOMZ also demonstrates improved performance compared to its base counterpart in most cases, except when classifying intents within the MASSIVE dataset. However, the Falcon models present a unique scenario where Falcon-Instruct, in contrast to other instruct models, does not consistently outperform the non-instruct base version. From Figure 3, it is evident that the MASSIVE dataset is the most challenging to solve. This variation may stem from the complexity of the dataset, posing challenges for the instruct model to generalize effectively to intricate patterns and implicit dependencies. This points to the need for further investigations and improved hyper-parameter tuning of instruction-tuned models.

Varying Model Size. We assess the impact of varying the model size for ILLUMINER instantiated with FLAN-T5_{LoRA}, and report the results in Figure 4. While there are notable gains with increased model size, especially for SF on all datasets and for IC on MASSIVE, smaller models perform nearly as well as the largest one for IC on SNIPS and MultiWoZ. This indicates that larger models excel in tasks with a vast set of labels. However, we observe a diminishing trend in performance gains after 3B, suggesting that leveraging models larger than 11B in the PEFT setting likely offers no advantages.

Varying Number of Examples per Label. In Figure 5, increasing the number of examples (k) per label for ILLUMINER with FLAN-T5-xxl_{LoRA}




Figure 6: Performance of FLAN-T5-xxl with different PEFT techniques. Bars indicate % of model parameters trained during PEFT.




Figure 7: Generalization of FLAN-T5-xxl with different PEFT techniques. Bars indicate average model performance across datasets.

shows minimal performance gains, except for SF on SNIPS and MASSIVE, where using all training instances leads to 5.2 and 7.7 percentage points improvement. This demonstrates that Instruct-LLMs in the PEFT setting are able to generalize effectively even with extremely limited fine-tuning data.

Different PEFT techniques. As illustrated in Figure 6, FLAN-T5-xxl was fine-tuned with different PEFT techniques to compare the trends. All the techniques were trained with parameters amounting to less than 0.1% of the model parameters. For the relatively easier IC task, the four techniques offer comparable performance across datasets. However, LoRA stands out as the most suitable for SF across all datasets. Notably, prefix- and prompt-tuned models exhibit poor SF performance in SNIPS and MASSIVE but perform similarly to LoRA and (IA)³ in MultiWoZ.

Generalization across datasets. We extended our evaluation across datasets for the models trained with diverse PEFT techniques to learn their cross-dataset generalization capabilities. For instance, LoRA adapters for FLAN-T5-xxl were trained individually on SNIPS for IC and SF, and evaluated on MASSIVE and MultiWoZ. Similar evaluations were conducted for all considered PEFT techniques and across all combinations of train/eval datasets. In Figure 7, we report generalization trends with average Accuracy for IC and average F1 for SF in the cross-dataset evaluation. It is evident that LoRA and (IA)³ offer impressive generalization over cross-dataset evaluation in both tasks. While prompt tuning shows comparable IC gener-

	SNIPS		MASSIVE		MultiWoZ	
	TEST−	TEST+	TEST−	TEST+	TEST−	TEST+
Intent Classification (Acc.)						
TRAIN−	0.930	0.946	0.809	0.777	0.973	0.986
TRAIN+	0.136	0.962	0.062	0.825	0.124	0.979
Slot Filling (F1)						
TRAIN−	0.002	0.143	0.033	0.207	0.161	0.389
TRAIN+	0.000	0.909	0.001	0.735	0.000	0.945

Table 6: Effect of fine-tuning and evaluation with (+) and without (−) labels in the instructions.

Model	MASSIVE									
	en		de		fr		it		es	
	IC	SF	IC	SF	IC	SF	IC	SF	IC	SF
flan-t5-xxl _{LoRA}	.833	.735	.783	.669	.770	.614	.745	.573	.733	.560
Multilingual LLMs										
mt5-xxl _{LoRA}	.795	.674	.775	.653	.784	.671	.773	.663	.773	.649
mt0-xxl _{LoRA}	.804	.689	.778	.662	.784	.660	.797	.658	.784	.631
mt0-xxl-mt _{LoRA}	.814	.700	.781	.689	.685	.679	.751	.679	.768	.670

Table 7: Performance on languages other than English in terms of intent accuracy (IC) and slot filling F1 (SF).

alization, it exhibits significantly lower performance in SF generalization.

Exposure of labels in instructions. To study the impact of the inclusion of labels (L and S for IC and SF, respectively, as defined in § 3) in instruction tuning, we conducted additional fine-tuning and evaluation of the FLAN-T5-xxl model without exposing these labels in the instruction. As shown in Table 6, models fine-tuned with instructions containing L and S (TRAIN+) outperform those without these labels during fine-tuning (TRAIN−). This difference is especially pronounced in SF, where it exceeds 62% on average. Interestingly, models trained with L and S perform relatively poorly when these labels are excluded during evaluation (TEST−). We conclude that fine-tuning models for IC and SF with labels in the instruction enables the generation of output labels from candidates L and S , offering improved generalization across domains and datasets compared to learning solely from model weights and data distributions.

Multilinguality. To investigate the applicability of our ILLUMINER framework for languages beyond English, we performed IC and SF experiments across five language splits within the MASSIVE dataset: English (en), German (de), French (fr), Italian (it), and Spanish (es). In Table 7, we report the performance of ILLUMINER instantiated with the following LoRA fine-tuned models:

- *FLAN-T5-xxl*, trained with mostly English texts.
- *mT5-xxl* (Xue et al., 2021), a multilingual variant of T5 covering 101 languages.
- *mT0-xxl(-mt)*, mT5 fine-tuned on a cross-lingual instruction dataset, xP3 (Muennighoff et al.,

2022). The *-mt* variant is recommended for prompting in non-English.

Multilingual LLMs generally exhibit lower performance on the English split compared to FLAN-T5. However, apart from IC on the German split, we observe the advantages of utilizing multilingual LLMs (mt5-xxl_{LoRA} and mt0-xxl_{LoRA}) for non-English input utterances, even when the task instructions and label descriptions are still in English. Instruction-tuned mT5, referred to as mT0, demonstrates superior performance across all considered languages, validating our previous observation that applying PEFT on Instruct-LLMs yields greater benefits. When employing mt0-xxl-mt_{LoRA}, we translated both task instructions and label descriptions into the respective languages of the input utterances, during both fine-tuning and inference stages. While we observe performance increase in SF when the prompts were translated, the same improvement was not always evident for IC. We conjecture that this discrepancy arises from translations often yielding longer and more ambiguous label descriptions, particularly noticeable for French.

7. Discussion

Based on the experimental outcomes, we record a few observations as shown in Table 8, discussing the shortcomings and advantages of few-shot learning and instruction tuning for IC and SF.

Ambiguous User Utterances. In TOD systems, models often deal with ambiguous user utterances as input, facing challenges in accurately identifying potential intents. Example 1 illustrates such an utterance, where the annotated intent is related to ‘pink’ as the smart lighting’s color, while ILLUMINER misunderstands it as a singer and GPT3.5 falls back to the out-of-scope intent (*be quirky*). Users convey intents in numerous ways, posing difficulties for models to generalize across variations.

Entity Disambiguation. In many cases including Example 1 and 2, words and phrases may refer to different entity types, requiring models to disambiguate them in the given context. However, given only few samples for either few-shot learning or fine-tuning, it is often hard for LLMs to understand patterns or guidelines employed by human annotators on deciding slot labels (e.g., ‘australian’ *time-zone* against the place ‘australia’).

Missing Context. Single-turn IC and SF is highly challenging due to limited context as compared to a multi-turn setting with previous turns in the conversation as context. In Example 3, context absence hinders the models to predict the expected intent.

Highly Correlated Labels. Highly correlated labels where the distinctions are often subtle and context-dependent, such as *entity-name* and *artist*,

ID	Problem Category	User utterance	Expected Label(s)	LLM Response	
				ILLUMINER	GPT3.5 <small>few-shot</small>
1	Ambiguous User Utterances	"pink is all we need"	<i>l</i> : change light color	express music liking	be quirky
2	Entity Disambiguation	"what's the time in australia"	<i>s</i> : place-name: australia	time-zone: australia	time-zone: australia
3	Missing Context	"remind me to do something then"	<i>l</i> : set a calendar event	set an alarm	set an alarm
4	Highly correlated labels	"put lindsey cardinale into my hillary clinton s women s history month playlist"	<i>s</i> : artist: lindsey cardinale	artist: lindsey cardinale	entity-name: lindsey cardinale
5	Hallucinations	"turn my morning alarm on"	<i>l</i> : set an alarm	turn an alarm on	set an alarm
6		"play it again please"	<i>s</i> : \emptyset	\emptyset	player-setting: repeat

Table 8: Problem categories with exemplars. *l* and *s* denote expected intents and slots. We report LLM predictions by ILLUMINER (flan-t5-xxl_{LoRA}) and GPT3.5_{few-shot}. Erroneous predictions are marked red.

making it challenging to precisely predict intents and slots given user utterances. This also points towards data inconsistencies and label noise in large datasets. Nevertheless, ILLUMINER correctly identified ‘*lindsey cardinale*’ in Example 4 as *artist*, supporting the hypothesis that fine-tuning may resolve such problems for most examples, if not entirely.

Hallucinations. LLMs are prone to hallucinations, as evidenced in our use case where they generate intents and slots absent in candidate labels or user utterances. Example 5 and 6 depict such a *factual mirrage* (Rawte et al., 2023) for IC and SF, respectively, where ILLUMINER generated the *turn an alarm on* intent not present in the candidate labels, and GPT3.5 generated the *player-setting: ‘repeat’* slot when ‘*repeat*’ is never mentioned. Approximately 2.94% of false positives for IC and 3.76% for SF, with ILLUMINER, fall into this error category.

For future research, we plan to extend our study to multi-turn settings to tackle context deficiency. Techniques like semantic-driven label mapping, confidence scoring for prediction reliability assessment, and requesting clarification could mitigate hallucination risks.

8. Conclusion

We introduced ILLUMINER for intent classification (IC) and slot filling (SF) with Instruct-LLMs. Our LoRA fine-tuned models surpass GPT3.5 in zero- and few-shot settings, as well as the state-of-the-art joint IC+SF approach. Notably, we achieve impressive results using less than 6% of the training data across benchmarks like SNIPS, MASSIVE and MultiWoZ. These findings have direct practical applications in task-oriented dialogue systems, enabling enhanced performance with reduced computational power and data annotation efforts.

Acknowledgement

This research was funded by the German Federal Ministry for Economic Affairs and Climate Action (BMWK) through the project OpenGPT-X (project no. 68GX21007D).

9. Bibliographical References

- Ebtesam Almazrouei et al. 2023. Falcon-40B: an open large language model with state-of-the-art performance.
- Tom Brown et al. 2020. [Language Models are Few-Shot Learners](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.
- Kai-Wei Chang, Ming-Hsin Chen, Yun-Ping Lin, Jing Neng Hsu, Paul Kuo-Ming Huang, Chien yu Huang, Shang-Wen Li, and Hung yi Lee. 2023. [Prompting and adapter tuning for self-supervised encoder-decoder speech model](#).
- Qian Chen, Zhu Zhuo, and Wen Wang. 2019. [BERT for Joint Intent Classification and Slot Filling](#). ArXiv:1902.10909 [cs].
- Leyang Cui, Yu Wu, Jian Liu, Sen Yang, and Yue Zhang. 2021. [Template-Based Named Entity Recognition Using BART](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 1835–1845, 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.
- Gabor Fuisz, Ivan Vulić, Samuel Gibbons, Inigo Casanueva, and Paweł Budzianowski. 2022. [Improved and efficient conversational slot labeling through question answering](#).

- Ian J. Goodfellow, Mehdi Mirza, Da Xiao, Aaron Courville, and Yoshua Bengio. 2015. [An Empirical Investigation of Catastrophic Forgetting in Gradient-Based Neural Networks](#). ArXiv:1312.6211 [cs, stat].
- Arshit Gupta, John Hewitt, and Katrin Kirchhoff. 2019. [Simple, Fast, Accurate Intent Classification and Slot Labeling for Goal-Oriented Dialogue Systems](#). In *Proceedings of the 20th Annual SIGdial Meeting on Discourse and Dialogue*, pages 46–55, Stockholm, Sweden. Association for Computational Linguistics.
- Raghav Gupta, Harrison Lee, Jeffrey Zhao, Yuan Cao, Abhinav Rastogi, and Yonghui Wu. 2022. [Show, Don't Tell: Demonstrations Outperform Descriptions for Schema-Guided Task-Oriented Dialogue](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4541–4549, Seattle, United States. Association for Computational Linguistics.
- Soyeon Caren Han, Siqu Long, Huichun Li, Henry Weld, and Josiah Poon. 2022. [Bi-directional Joint Neural Networks for Intent Classification and Slot Filling](#). ArXiv:2202.13079 [cs].
- Yutai Hou, Cheng Chen, Xianzhen Luo, Bohan Li, and Wanxiang Che. 2022. [Inverse is Better! Fast and Accurate Prompt for Few-shot Slot Tagging](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 637–647, Dublin, Ireland. Association for Computational Linguistics.
- Yutai Hou, Yongkui Lai, Cheng Chen, Wanxiang Che, and Ting Liu. 2021. [Learning to Bridge Metric Spaces: Few-shot Joint Learning of Intent Detection and Slot Filling](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 3190–3200, Online. Association for Computational Linguistics.
- Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin De Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. 2019. [Parameter-Efficient Transfer Learning for NLP](#). In *Proceedings of the 36th International Conference on Machine Learning*, pages 2790–2799. PMLR. ISSN: 2640-3498.
- Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, Lu Wang, and Weizhu Chen. 2022a. [LoRA: Low-Rank Adaptation of Large Language Models](#). In *International Conference on Learning Representations*.
- Yushi Hu, Chia-Hsuan Lee, Tianbao Xie, Tao Yu, Noah A. Smith, and Mari Ostendorf. 2022b. [In-Context Learning for Few-Shot Dialogue State Tracking](#). In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 2627–2643, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Vojtěch Hudeček and Ondřej Dušek. 2023. [Are LLMs All You Need for Task-Oriented Dialogue?](#) ArXiv:2304.06556 [cs].
- Chia-Chien Hung, Anne Lauscher, Simone Paolo Ponzetto, and Goran Glavaš. 2022. [Ds-tod: Efficient domain specialization for task oriented dialog](#).
- Jason Krone, Yi Zhang, and Mona Diab. 2020. [Learning to Classify Intents and Slot Labels Given a Handful of Examples](#). In *Proceedings of the 2nd Workshop on Natural Language Processing for Conversational AI*, pages 96–108, Online. Association for Computational Linguistics.
- Sang Yun Kwon, Gagan Bhatia, Elmoatez Billah Nagoudi, Alcides Alcoba Inciarte, and Muhammad Abdul-mageed. 2023. [SIDLR: Slot and intent detection models for low-resource language varieties](#). In *Tenth Workshop on NLP for Similar Languages, Varieties and Dialects (VarDial 2023)*, pages 241–250, Dubrovnik, Croatia. Association for Computational Linguistics.
- Brian Lester, Rami Al-Rfou, and Noah Constant. 2021. [The Power of Scale for Parameter-Efficient Prompt Tuning](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 3045–3059, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Moxin Li, Wenjie Wang, Fuli Feng, Jizhi Zhang, and Tat-Seng Chua. 2023a. [Robust instruction optimization for large language models with distribution shifts](#).
- Xiang Lisa Li and Percy Liang. 2021. [Prefix-Tuning: Optimizing Continuous Prompts for Generation](#). In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 4582–4597, Online. Association for Computational Linguistics.
- Xuefeng Li, Liwen Wang, Guanting Dong, Keqing He, Jinzheng Zhao, Hao Lei, Jiachi Liu, and Weiran Xu. 2023b. [Generative Zero-Shot Prompt Learning for Cross-Domain Slot Filling with Inverse Prompting](#). In *Findings of the Association for Computational Linguistics: ACL 2023*, pages

- 825–834, Toronto, Canada. Association for Computational Linguistics.
- Yen-Ting Lin, Alexandros Papangelis, Seokhwan Kim, Sungjin Lee, Devamanyu Hazarika, Mahdi Namazifar, Di Jin, Yang Liu, and Dilek Hakkani-Tur. 2023. [Selective In-Context Data Augmentation for Intent Detection using Pointwise V-Information](#). In *Proceedings of the 17th Conference of the European Chapter of the Association for Computational Linguistics*, pages 1463–1476, Dubrovnik, Croatia. Association for Computational Linguistics.
- Haokun Liu, Derek Tam, Mohammed Muqeeth, Jay Mohta, Tenghao Huang, Mohit Bansal, and Colin Raffel. 2022a. [Few-shot parameter-efficient fine-tuning is better and cheaper than in-context learning](#).
- Haokun Liu, Derek Tam, Mohammed Muqeeth, Jay Mohta, Tenghao Huang, Mohit Bansal, and Colin A. Raffel. 2022b. [Few-Shot Parameter-Efficient Fine-Tuning is Better and Cheaper than In-Context Learning](#). *Advances in Neural Information Processing Systems*, 35:1950–1965.
- 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](#). ArXiv:1907.11692 [cs].
- Jianqiang Ma, Zeyu Yan, Chang Li, and Yang Zhang. 2021. [Frustratingly Simple Few-Shot Slot Tagging](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 1028–1033, Online. Association for Computational Linguistics.
- Niklas Muennighoff, Thomas Wang, Lintang Sutawika, Adam Roberts, Stella Biderman, Teven Le Scao, M Saiful Bari, Sheng Shen, Zheng Xin Yong, Hailey Schoelkopf, Xiangru Tang, Dragomir Radev, Alham Fikri Aji, Khalid Al-mubarak, Samuel Albanie, Zaid Alyafeai, Albert Webson, Edward Raff, and Colin Raffel. 2023. [Crosslingual Generalization through Multitask Finetuning](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 15991–16111, Toronto, Canada. Association for Computational Linguistics.
- Long Ouyang, Jeffrey Wu, Xu Jiang, Diogo Almeida, Carroll Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul F. Christiano, Jan Leike, and Ryan Lowe. 2022. [Training language models to follow instructions with human feedback](#). *Advances in Neural Information Processing Systems*, 35:27730–27744.
- Soham Parikh, Mitul Tiwari, Prashil Tumbade, and Quaizar Vohra. 2023. [Exploring Zero and Few-shot Techniques for Intent Classification](#). In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 5: Industry Track)*, pages 744–751, Toronto, Canada. Association for Computational Linguistics.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. [Language Models are Unsupervised Multitask Learners](#). Technical report, OpenAI.
- Vipula Rawte, Swagata Chakraborty, Agnib Pathak, Anubhav Sarkar, S. M. Towhidul Islam Tonmoy, Aman Chadha, Amit P. Sheth, and Amitava Das. 2023. [The Troubling Emergence of Hallucination in Large Language Models – An Extensive Definition, Quantification, and Prescriptive Remediations](#). ArXiv:2310.04988 [cs].
- Andy Rosenbaum, Saleh Soltan, Wael Hamza, Yannick Versley, and Markus Boese. 2022. [LINGUIST: Language Model Instruction Tuning to Generate Annotated Utterances for Intent Classification and Slot Tagging](#). In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 218–241, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.
- Timo Schick and Hinrich Schütze. 2021. [Exploiting Cloze-Questions for Few-Shot Text Classification and Natural Language Inference](#). In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 255–269, Online. Association for Computational Linguistics.
- Yixuan Su, Lei Shu, Elman Mansimov, Arshit Gupta, Deng Cai, Yi-An Lai, and Yi Zhang. 2022. [Multi-Task Pre-Training for Plug-and-Play Task-Oriented Dialogue System](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 4661–4676, Dublin, Ireland. Association for Computational Linguistics.
- Hugo Touvron et al. 2023a. [Llama 2: Open Foundation and Fine-Tuned Chat Models](#). ArXiv:2307.09288 [cs].
- Hugo Touvron et al. 2023b. [LLaMA: Open and Efficient Foundation Language Models](#). ArXiv:2302.13971 [cs].

- Lewis Tunstall, Nils Reimers, Unso Eun Seo Jo, Luke Bates, Daniel Korat, Moshe Wasserblat, and Oren Pereg. 2022. [Efficient Few-Shot Learning Without Prompts](#).
- Qingyue Wang, Yanan Cao, Piji Li, Yanhe Fu, Zheng Lin, and Li Guo. 2022a. [Slot Dependency Modeling for Zero-Shot Cross-Domain Dialogue State Tracking](#). In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 510–520, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.
- Weizhi Wang, Zhirui Zhang, Junliang Guo, Yinpei Dai, Boxing Chen, and Weihua Luo. 2022b. [Task-oriented dialogue system as natural language generation](#). In *Proceedings of the 45th International ACM SIGIR Conference on Research and Development in Information Retrieval*. ACM.
- Jason Wei, Maarten Bosma, Vincent Y. Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M. Dai, and Quoc V. Le. 2022. [Fine-tuned Language Models Are Zero-Shot Learners](#). ArXiv:2109.01652 [cs].
- Henry Weld, Xiaoqi Huang, Siqu Long, Josiah Poon, and Soyeon Caren Han. 2023. [A Survey of Joint Intent Detection and Slot Filling Models in Natural Language Understanding](#). *ACM Computing Surveys*, 55(8):1–38.
- BigScience Workshop et al. 2022. [BLOOM: A 176B-Parameter Open-Access Multilingual Language Model](#).
- Tianbao Xie, Chen Henry Wu, Peng Shi, Ruiqi Zhong, Torsten Scholak, Michihiro Yasunaga, Chien-Sheng Wu, Ming Zhong, Pengcheng Yin, Sida I. Wang, Victor Zhong, Bailin Wang, Chengzu Li, Connor Boyle, Ansong Ni, Ziyu Yao, Dragomir Radev, Caiming Xiong, Lingpeng Kong, Rui Zhang, Noah A. Smith, Luke Zettlemoyer, and Tao Yu. 2022. [UnifiedSKG: Unifying and Multi-Tasking Structured Knowledge Grounding with Text-to-Text Language Models](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 602–631, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Linting Xue, Noah Constant, Adam Roberts, Mihir Kale, Rami Al-Rfou, Aditya Siddhant, Aditya Barua, and Colin Raffel. 2021. [mT5: A massively multilingual pre-trained text-to-text transformer](#).
- Dian Yu, Luheng He, Yuan Zhang, Xinya Du, Panupong Pasupat, and Qi Li. 2021. [Few-shot Intent Classification and Slot Filling with Retrieved Examples](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 734–749, Online. Association for Computational Linguistics.
- Jianguo Zhang, Trung Bui, Seunghyun Yoon, Xiang Chen, Zhiwei Liu, Congying Xia, Quan Hung Tran, Walter Chang, and Philip Yu. 2021. [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.
- Shengyu Zhang, Linfeng Dong, Xiaoya Li, Sen Zhang, Xiaofei Sun, Shuhe Wang, Jiwei Li, Runyi Hu, Tianwei Zhang, Fei Wu, and Guoyin Wang. 2023. [Instruction Tuning for Large Language Models: A Survey](#). ArXiv:2308.10792 [cs].
- Lianmin Zheng, Wei-Lin Chiang, Ying Sheng, Siyuan Zhuang, Zhanghao Wu, Yonghao Zhuang, Zi Lin, Zhuohan Li, Dacheng Li, Eric P. Xing, Hao Zhang, Joseph E. Gonzalez, and Ion Stoica. 2023. [Judging LLM-as-a-judge with MT-Bench and Chatbot Arena](#).

10. Language Resource References

- Budzianowski, Paweł and Wen, Tsung-Hsien and Tseng, Bo-Hsiang and Casanueva, Inigo and Ultes, Stefan and Ramadan, Osman and Gašić, Milica. 2018. *Multiwoz—a large-scale multi-domain wizard-of-oz dataset for task-oriented dialogue modelling*.
- Coucke, Alice and Saade, Alaa and Ball, Adrien and Bluche, Théodore and Caulier, Alexandre and Leroy, David and Doumouro, Clément and Gisselbrecht, Thibault and Caltagirone, Francesco and Lavril, Thibaut and others. 2018. *Snips voice platform: an embedded spoken language understanding system for private-by-design voice interfaces*.
- FitzGerald, Jack and Hench, Christopher and Peris, Charith and Mackie, Scott and Rottmann, Kay and Sanchez, Ana and Nash, Aaron and Urbach, Liam and Kakarala, Vishesh and Singh, Richa and others. 2022. *Massive: A 1m-example multilingual natural language understanding dataset with 51 typologically-diverse languages*.
- Longpre, Shayne and Hou, Le and Vu, Tu and Webson, Albert and Chung, Hyung Won and Tay, Yi and Zhou, Denny and Le, Quoc V. and Zoph, Barret and Wei, Jason and Roberts, Adam. 2023.

The Flan Collection: Designing Data and Methods for Effective Instruction Tuning. [\[link\]](#).

Muennighoff, Niklas and Wang, Thomas and Sutawika, Lintang and Roberts, Adam and Biderman, Stella and Scao, Teven Le and Bari, M Saiful and Shen, Sheng and Yong, Zheng-Xin and Schoelkopf, Hailey and others. 2022. *xP3: Crosslingual generalization through multi-task finetuning.* BigScience, distributed via Hugging Face. [\[link\]](#).

Penedo, Guilherme and Malartic, Quentin and Hesslow, Daniel and Cojocaru, Ruxandra and Cappelli, Alessandro and Alobeidli, Hamza and Pannier, Baptiste and Almazrouei, Ebtesam and Launay, Julien. 2023. *The RefinedWeb Dataset for Falcon LLM: Outperforming Curated Corpora with Web Data, and Web Data Only.* arXiv. [\[link\]](#).

Xu, Can and Sun, Qingfeng and Zheng, Kai and Geng, Xiubo and Zhao, Pu and Feng, Jiazhan and Tao, Chongyang and Jiang, Daxin. 2023a. *Wizardlm: Empowering large language models to follow complex instructions.* WizardLM, distributed via Hugging Face. [\[link\]](#).

Xu, Canwen and Guo, Daya and Duan, Nan and McAuley, Julian. 2023b. *Baize: An Open-Source Chat Model with Parameter-Efficient Tuning on Self-Chat Data.* arXiv. [\[link\]](#).