

# StructLM: Towards Building Generalist Models for Structured Knowledge Grounding

Alex Zhuang<sup>1,\*</sup>, Ge Zhang<sup>1,2,7\*</sup>,  
Tianyu Zheng<sup>2</sup>, Xinrun Du<sup>2</sup>, Junjie Wang<sup>3,4</sup>, Weiming Ren<sup>1,7</sup>,  
Stephen W. Huang<sup>6</sup>, Jie Fu<sup>2,4†</sup>, Xiang Yue<sup>2,5</sup>, Wenhui Chen<sup>1,2,7†</sup>

<sup>1</sup>University of Waterloo, <sup>2</sup>Multimodal Art Projection Research Community,

<sup>3</sup>Waseda University, <sup>4</sup>HKUST, <sup>5</sup>Ohio State University, <sup>6</sup>harmony.ai, <sup>7</sup>Vector Institute  
<https://tiger-ai-lab.github.io/StructLM/>

## Abstract

Structured data sources, such as tables, graphs, and databases, are ubiquitous knowledge sources. Despite the demonstrated capabilities of large language models (LLMs) on plain text, their proficiency in interpreting and utilizing structured data remains limited. Our investigation reveals a notable deficiency in LLMs’ ability to process structured data, e.g., ChatGPT lags behind state-of-the-art (SoTA) model by an average of 35%. To augment the Structured Knowledge Grounding (SKG) capabilities in LLMs, we have developed a comprehensive instruction tuning dataset comprising 1.1 million examples. Utilizing this dataset, we train a series of models, referred to as StructLM, based on Mistral and the CodeLlama model family, ranging from 7B to 34B parameters. Our StructLM series surpasses task-specific models (Xie et al., 2022) on 16 out of 18 evaluated datasets and establishes new SoTA performance on 8 SKG tasks. Furthermore, StructLM demonstrates strong generalization across 6 novel held-out SKG tasks, outperforming TableLlama by an average of 35% and Flan-UL2 20B by an average of 10%. Contrary to expectations, we observe that scaling model size offers marginal benefits, with StructLM-34B showing only slight improvements over StructLM-7B. This suggests that structured knowledge grounding is still a challenging task and requires more innovative design to push to a new level. We release the model weights and training dataset to the community, along with relevant code on Github.

## 1 Introduction

Traditionally, users need to write programs to interface with structured data like tables, databases, knowledge graphs, etc. This often requires that they master a the domain-specific language like SQL, SPARQL, etc. or develop domain specific skills for that structured type. Recently, researchers have explored the possibility of automating the interface with natural language to enable potential use cases in question-answering (Pasupat & Liang, 2015; Zhong et al., 2017; Nan et al., 2022), summarization (Parikh et al., 2020; Nan et al., 2021; Bao et al., 2018), and fact verification (Aly et al., 2021; Chen et al., 2019; Gupta et al., 2020b), among others, all grounded to a structured knowledge source. This effort can lower the barrier for end users to access a massive amount of structured data.

Previous work (Yu et al., 2020; Liu et al., 2021; Xie et al., 2022; Zhang et al., 2023) has been mostly focused on building task-specific models for different tasks with rather limited generalization ability. Building a generalist structure knowledge grounding (SKG) system across a wide range of tasks proves to be challenging. This is mainly due to the heterogeneity of data format and use cases. We evaluated GPT-3.5-Turbo (Jiang et al., 2023) on 18 SKG tasks and observed that its performance is on average 35% lower than the SoTA specialized models. It shows that the LLM’s ability on SKG is heavily overlooked during pre-training.

In this paper, we explore the possibility of building a generalist model based on LLMs that can ground on diverse types of structure and unstructured knowledge to interface with humans. Specifically, we construct a large data set of over a million instruction-following examples, a majority of which is SKG data, along with additional general instruction-following data, which we find improves generalizability. We fine-tune models at three scales: 7B, 13B, and 34B, based on Mistral and the CodeLlama family of code foundation models. When compared to USKG, we find that our models surpass these single-task models on 16 of 18 tasks. StructLM achieves SoTA on 8 of 18 evaluated tasks, beating ChatGPT by a huge margin. Moreover, we show that compared with TableLlama, a recent open language model finetuned for tabular inference tasks, our held-out performance is vastly superior. Compared with Flan-UL2-20B, our models also generalize better overall on the held out tasks by as much as 10% on average.

We study the performance of StructLM, namely whether the model experiences cross-task generalization benefits from the dataset mixture, and find that our multi-task model performs significantly better overall than single-task models of the exact same parameter scale. We also study the effect of different pretraining data on our finetuned performance to determine whether special pretraining regimes, such as code or math, contribute to effective SKG reasoning ability. We find that code pretraining is the most effective. We perform additional ablations to confirm our results and support our claims. Our contributions are:

- We construct a large SKG instruction-tuning dataset with 1.1 million samples. We train and release our models that outperform the previous 3B USKG fine-tuned on individual tasks on a total of 16 of 18 tasks. StructLM also achieves SoTA results on 8 of them.
- We show that StructLM is able to show strong zero-shot generalization capability on unseen structure knowledge grounding tasks, which was not shown by previous models. StructLM outperforms Flan-UL2 20B on 3 of 5 tasks and TableLlama on all 6 tasks. The average absolute improvement over them are 10% and 35%.
- We find that mixing in general-domain instruction-tuning data during finetuning preserves generalization ability, and that code-pretrained base models can improve model performance on the SKG tasks.

## 2 Related Works

### 2.1 Solving SKG tasks

Structured knowledge, such as web tables, knowledge graphs, and databases, have long been the subject of study in knowledge grounding. However, SKG tasks have heterogeneous data formats which have inspired methods that leverage specific training setups to




Figure 1: StructLM can ground on structured and unstructured knowledge to respond to human queries. The previous SoTA was attained by many different task-specific models like TAPEx (Liu et al., 2021), USKG (Xie et al., 2022), TableLlama (Zhang et al., 2023), BINDER-Codex (Cheng et al., 2022), etc. StructLM beats the SoTAs on seven SKG tasks.

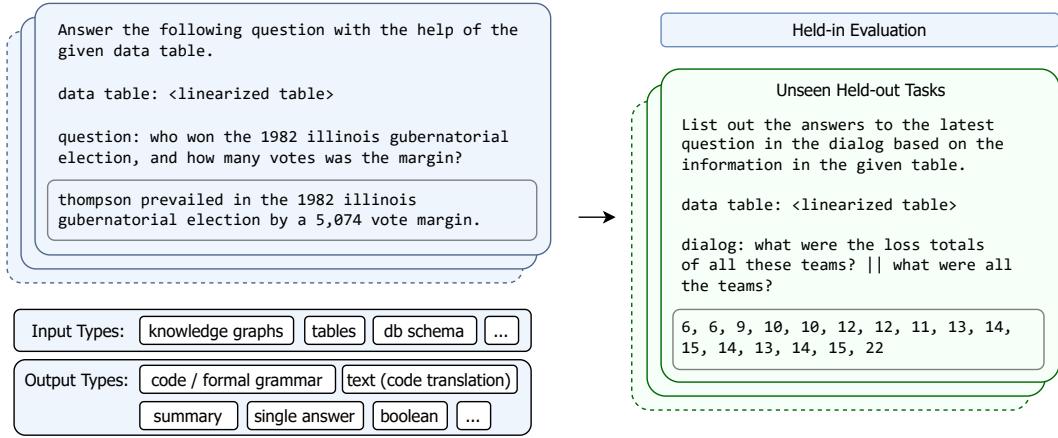


Figure 2: Overview of StructLM. This figure illustrates the prompting structure of StructLM, highlighting its capability to process various forms of structured data beyond linearized data tables, including linearized database schemas and knowledge graphs.

learn those representations. For example, PTab (Liu et al., 2022) and MultiHiertt (Zhao et al., 2022) learn the contextual representation of tabular data by incorporating semantic information through specific training methods or reasoning approaches. RASAT (Qi et al., 2022a) integrates relation-aware self-attention with the Transformer seq2seq architecture and utilizes various relational structures to address SQL problems. TAPEX (Liu et al., 2021) conducts pretraining over tabular/database data with the help of an SQL executor to provide supervision.

More recently, methods have begun to move away from these auxiliary task-specific structures. USKG (Xie et al., 2022) were the first to unify many SKG tasks into a sequence-to-sequence format, allowing them to be aggregated into the same data mixture. However, they were not able to show strong performance improvements to constructing a multi-task mix of SKG data over task-specific tuning methods. StructGPT (Jiang et al., 2023) represents a line of work that uses prompting frameworks on powerful LLMs to solve tasks with more robustness and accuracy. In contrast, our work examines open models and tries to assess their fundamental capabilities. Contemporary to our work, TableLlama (Zhang et al., 2023) has argued that tabular data deserves special attention. Focusing on this domain, their method fine-tunes on several new tabular tasks to improve table understanding, and operates on a longer 8k context length. These improvements can be additive to our work.

## 2.2 LLMs with Instruction Tuning

Instruction-tuning (IT) has been popularized as a method to address the gap between training objectives and user goals in LLMs. This technique involves additional training of LLMs using pairs of instructions and outputs. IT enhances both the controllability and the predictability of the models, aligning them more closely with user expectations. Furthermore, recent studies such as FLAN (Wei et al., 2022), UL2 (Tay et al., 2023a), and Llama2 (Touvron et al., 2023) have shown that IT can improve the performance of downstream tasks through multi-task learning across diverse data types. While FLAN-UL2 trains on a subset of 11 SKG tasks, it also trains on many more unrelated language tasks. In our work, by focusing on SKG data, we hope to provide a focused study that can act as a reference for future work to improve performance on this task type.

## 2.3 Reasoning Capability in LLMs

Reasoning stands as a pivotal skill for LLMs in the development of real-world AI applications which would enable the autonomous completion of many thought-intensive tasks viewed traditionally to require human thinking, like programming or mathematical problem-solving (Li et al., 2022). Recent studies (Li et al., 2022; 2023c; Rozière et al., 2023;




Figure 3: Breakdown of Structured Knowledge Types and Tasks in the Training Data. On the left side, we see a coarse breakdown of the different categories of structured inputs in our dataset. On the right side, we see an overview of the task groups that are represented for those structured knowledge types.

Azerbaiyev et al., 2023a) indicate that LLMs trained on code and mathematical datasets exhibit profound reasoning skills, and can even achieve performance on par with human levels. For example, CodeLlama (Rozière et al., 2023), a foundation model trained on more programming data, has significantly improved reasoning capabilities across a variety of programming and mathematical benchmarks. Furthermore, Llemma (Azerbaiyev et al., 2023a) continues to pretrain the CodeLlama model on a mix of scientific papers, math-related web data, and mathematical code. Its results show excellent reasoning capabilities on the MATH benchmark (Hendrycks et al., 2021) and the ability to prove theorems without further fine-tuning. On the fine-tuning side, WizardMath (Luo et al., 2023a), and WizardCoder (Luo et al., 2023c) have shown the effectiveness of instruction tuning on reasoning capabilities, given high quality data.

In this work, we view structured data as a third testbed for a different kind of reasoning within LLMs. We posit that in addition to mathematical or logical reasoning, the ability to recognize and make use of patterns within a structured input indicates that a model has robust representations of relationships in data. These representations may serve as a strong prior for further reasoning downstream.

### 3 Method

#### 3.1 Dataset Curation

Motivated by the goal of training a language model generally capable of a wide range of structured data tasks, we select a total of 25 SKG tasks to study. We report results on 18 held-in and 6 held-out tasks, where each held-out task meant to roughly evaluate the generalization capability of a held-in task group. In total, our held-in training dataset contains approximately 700k SKG examples. We describe the held-in dataset groups below.

**Data to Text Generation.** This group of datasets deals with the summarization or interpretation of structured data from tables to knowledge triples to formal languages. Their inclusion is motivated by the idea that useful LMs should be able to make sense of a wide variety of structured information and map it to meaning in natural language. The corresponding held-out dataset for this task group is intended to be WikiTableText.

**Table based Question Answering.** This group of datasets deals specifically with tabular data, optionally combined with text passages. LMs which are able to accurately answer questions and retrieve information from tables can be widely useful as assistants. The corresponding held-out dataset for this task group is SQA.

**Knowledge-grounded Conversations.** This group of tasks evaluates knowledge grounding in-conversation. Humans naturally interface with LMs through chat, and enabling this capability can lower the barrier to accessing the information in stored structured data. These tasks track user intention through provided dialogue and ask the model to provide an answer to the latest question. The held-out dataset for this task group is CoSQL.

**Fact verification.** One common use case for tables is to reference facts. In addition to question answering, the ability to reliably determine if data in a table supports a statement signals the existence of a robust representation of the table’s data. The held-out dataset for this task group is InfoTabs.

**SQL or domain-specific languages** SQL is the language most commonly used to interface with structured data today. Understanding how to write SQL also requires understanding of abstractions of tables and how they are linked together. In other domain-specific languages, the MTOP task measures a model’s ability to parse a specification and generate an API call, which sees potential in LLM tool use (e.g., (Qin et al., 2023)). The corresponding held-out dataset for this task group is intended to be BIRD (Li et al., 2023b), which further tests SQL generation abilities.

**Mathematical reasoning.** An analysis of tabular data may also require performing quick mathematical computations over their contents. Performance on these datasets tells us how well models can combine both structured knowledge and mathematical reasoning. As there are currently a limited number of datasets that combine mathematical reasoning with SKG, this category includes just TabMWP in the held-in corpus. We set FinQA as a challenging held-out dataset analogue. Not only does it require financial domain knowledge, but it combines tabular information with long text passages, and requires the generation of mathematical code.

**General instruction data.** In addition to the SKG datasets within the held-in dataset mixture, we also included general instruction tuning data without any structured knowledge component, to maintain the instruction-following ability of our model. We use SlimOrca (Lian et al., 2023), which is constructed from cleaned GPT-4 responses to a number of prompts from existing general large-scale instruction-following datasets. We detect no signs of data contamination for our held-out datasets based on our ablation results. We give a detailed overview of all dataset statistics in Table 1.

### 3.2 Instruction Finetuning Approach

To instruction tune our model, each example in our dataset consists of a system prompt, instruction, input, and output. For all SKG data examples, we use the same system prompt. For each dataset, we write 10 instruction variations, which are randomized when constructing the training samples. For SKG data, the input is composed of a combination of a structured knowledge input and accompanying text that could be a question, statement, or anything that would be required to specify the task. The prompt is provided in Figure 6.

### 3.3 Training and Evaluation Details

The base models for StructLM are the CodeLlama-Instruct family of models (Rozière et al., 2023). We finetune all models with a batch size of 512 for 3 epochs on A800 gpus. We train our 7-34B models on 16-64 GPUs using DeepSpeed ZeRO-3 (Rasley et al., 2020) such that the total training time for each model is between 3-5 days. This training setup is largely in line with community conventions, such as the settings used for the WizardLM (Xu et al., 2023), WizardMath (Luo et al., 2023a), and WizardCoder (Luo et al., 2023c) models.

We follow the structured data linearization conventions in USKG (Xie et al., 2022). However, we use a different truncation scheme as described below. During training, we maintain a maximum sequence length of 2048. To preserve as much input context as possible, when truncating we consider the combined token length of the prompt input and output label. We truncate only the structured knowledge portion of the input so that the example becomes at most 2048 tokens long. As shown in the dataset statistics in Table 1, setting the max token length of the examples in our dataset to 2048 allows nearly all examples to fit within the



Dataset	Overall Length		Train				Test			
	Input (avg)	Output (avg)	Count	Input (max)	Output (max)	# Trunc.	Count	Input (max)	Output (max)	# Trunc.
TabMWP	207.8	4.5	23059	709	33	0	7686	703	31	0
ToTTo	251.8	31.0	120761	2040	155	467	7700	2048	119	31
GrailQA	281.0	44.1	44337	884	134	0	6463	546	123	0
SQL2Text	122.3	18.1	5600	337	61	0	1034	245	38	0
MMQA	656.2	7.7	15688	2047	146	234	1501	2048	94	11
Spider	266.6	36.0	7000	1369	226	0	1034	453	146	0
KVRet	573.4	17.1	6288	1217	161	0	807	1147	82	0
HybridQA	700.4	6.8	62682	2047	91	200	3466	2048	79	6
SParC	276.3	32.6	12059	1417	226	0	1625	467	146	0
CompWebQ	1350.3	11.9	27639	2047	321	321	2816	2048	256	8
TabFact	660.1	4.6	92283	2045	5	2	12779	1687	4	0
WikiTQ	831.8	5.8	11321	2028	273	0	4344	2048	148	10
WikiSQL	689.2	7.1	56355	2047	518	16	15878	2048	244	1
FeTaQA	653.2	38.8	7326	1853	158	0	2003	1548	114	0
FEVEROUS	799.3	3.4	40669	2047	5	2052	4285	2048	4	195
MultiWOZ	777.2	154.5	56668	1656	196	0	7368	1344	185	0
DART	133.7	30.3	62659	406	258	0	5097	261	109	0
Logic2Text	166.1	26.9	8566	358	67	0	1092	347	60	0
MTOP	961.0	34.4	15667	1002	215	0	4386	990	113	0
SlimOrca	278.9	152.4	512069	2047	1808	0	-	-	-	-
BIRD	439.8	63.3	9428	1992	347	99	1534	1214	386	0
CoSQL	287.4	34.9	9502	1640	226	0	1300	535	190	0
SQA	656.9	34.9	12275	1812	1012	2	3011	1725	769	0
Infotabs	276.9	3.7	16538	1009	5	0	5400	1105	4	0
WikiTableText	149.6	27.4	10000	313	97	0	2000	226	89	0
Finqa	1230.3	21.0	6251	2040	72	186	1147	2048	61	25

Table 1: Token sequence length statistics for each dataset in our train and test sets. Input and output statistics are in tokens. We report the number of examples which have been truncated in each dataset.

context window with rare truncations. We discard examples for which even this structured input truncation is insufficient (e.g. the output is too long). During inference, we set the input token length to 2048, to allow even more structured information to be placed within the input context. We set the maximum generation length to 1024, which is sufficient for all correct responses in all datasets. For each model, including our single-task finetuned models, we choose the best performing checkpoint of the 3-epoch checkpoints.

## 4 Experiments

**Baselines** Firstly, to illustrate the current performance of language models on SKG tasks, we evaluate ChatGPT (GPT-3.5-turbo) and the base model CodeLlama-7B-Instruct under a 1-shot setting. Our prompting scheme, using the same linearized knowledge structures as in our held-in training, sees them struggle across the board with many tasks due to the unseen structure knowledge format. Although ChatGPT is superior on text-based tasks, its performance is lackluster on SKG tasks. Its gap with SoTA models is as significant as 35%.

**Held-in Results** To evaluate the benefits of our instruction-tuning dataset mix, we also run single-task baseline (each a 7B model) on each task and report their individual performance. We again use CodeLlama-7B-Instruct as the base model for each, and match each single task model on the same number of epochs (3) that was used to train the multitask models, ensuring that each model has seen the same data the same number of times. We observe that our multi-task models outperform these single-task models on nearly every task, with some by a considerable margin of up to 7%. This demonstrates the effectiveness of our instruction tuning dataset and supports the presence of cross-task generalization.

When compared to the 18 task-specific USKG models, StructLM-7B can surpass USKG by a average of 2%. From a parameter-count perspective, each of the USKG models is a T5-3B model, which means over the entire held-in set, these results require 54B parameters. Our 7B-M StructLM in comparison can be viewed as being nearly 8x as parameter efficient while

Dataset	Metric	SoTA	ChatGPT	Base-M	Base	ST	FLAN-UL2	TableLlama	USKG	StructLM (Ours)				$\Delta$
Size	-	-	-	7B	7B	7B×18	20B	7B	3B×18	7B-M	7B	13B	34B	-
Held In														
ToTo	BLEU	49.9	20.7	17.9	17.5	48.8	-	-	49.0	49.8	49.4	49.3	<b>50.2</b>	+0.3
GraillQA	EM	77.1	9.3	1.5	1.0	77.0	-	-	70.1	81.2	80.4	79.2	<b>82.2</b>	+5.1
SQL2Text	Blec	94.8	88.6	90.7	82.9	95.2	-	-	94.8	<b>95.2</b>	93.8	88.5	92.6	+0.4
MMQA	F1	85.3	59.6	41.5	30.7	81.5	-	-	85.3	85.5	85.2	86.0	<b>88.1</b>	+2.8
Spider	EM	<b>80.5</b>	43.8	31.0	5.2	67.3	-	-	71.8	72.4	72.4	74.1	<b>74.6</b>	-5.9
KVRet	All Micro	67.9	52.9	34.4	39.5	70.9	-	-	67.9	72.2	<b>72.6</b>	69.5	69.3	+4.7
HybridQA	Acc	<b>68.4</b>	23.7	12.9	2.3	58.4	61.0	-	59.4	<b>62.6</b>	59.2	59.1	61.1	-5.8
SParC	EM	<b>68.2</b>	32.2	23.7	3.2	62.3	-	-	61.5	63.3	61.9	<b>64.9</b>	63.4	-3.3
CompWebQ	Acc	76.8	48.9	30.9	3.1	75.6	75.9	-	73.3	79.9	78.3	80.4	<b>81.9</b>	+5.1
TabFact	Acc	<b>93.0</b>	62.4	25.7	0.0	79.6	87.1	82.5	83.7	84.6	80.8	84.7	<b>86.6</b>	-6.4
WikiTQ	All Ex	<b>65.9</b>	24.8	6.7	0.2	45.7	54.6	-	49.3	<b>56.8</b>	50.1	53.4	55.7	-9.1
WikiSQL	All Ex	<b>93.0</b>	31.5	21.5	0.4	86.5	87.3	-	86.0	87.0	<b>88.7</b>	87.2	87.6	-4.3
FeTaQA	BLEU	<b>39.0</b>	7.4	13.7	5.6	33.8	35.8	39.0	33.4	37.5	36.0	35.6	<b>37.5</b>	-1.5
FEVEROUS	Acc	85.6	57.8	73.2	58.4	78.1	85.6	-	82.4	<b>85.9</b>	84.4	85.0	85.7	+0.3
MultiWOZ	Joint Acc	<b>60.6</b>	8.9	0.3	0.0	53.0	-	-	55.4	<b>55.4</b>	54.5	53.0	53.8	-5.2
DART	BLEU	52.0	59.0	47.4	54.6	60.3	50.4	-	46.7	<b>63.2</b>	62.2	61.4	61.8	+11.2
Logic2Text	Blec	<b>95.3</b>	78.5	81.5	59.1	89.5	-	-	91.4	89.5	88.9	<b>90.1</b>	89.1	-5.2
MTOP	EM	<b>87.5</b>	1.4	0.8	0.0	77.4	87.5	-	86.8	75.8	81.2	81.6	<b>82.1</b>	-5.4
<b>Average</b>		74.9	39.5	30.8	20.2	68.2	-	-	69.3	72.1	71.1	71.3	<b>72.6</b>	-1.2
Held Out														
BIRD	Acc	36.6*	21.8	11.5	0.0	24.4*	1.0	0	0	22.8	22.3	22.8	<b>24.7</b>	+2.9
CoSQL	EM	58.3*	33.7	26.5	0.2	52.4*	5.1	0	0	52.8	49.8	52.2	<b>55.0</b>	+21.3
SQA	Acc	70.5*	18.7	7.4	2.3	60.4*	70.1*	0	0	42.6	<b>49.7</b>	36.1	44.2	+31
Infotabs	Acc	75.6*	46.9	49.1	40.2	68.7*	<b>70.3</b>	16.2	0	47.2	55.3	58.1	<b>61.8</b>	-8.5
WikiTableText	BLEU	33.7*	3.8	3.9	5.7	39.8*	<b>19.4</b>	3.4	0	<b>17.1</b>	8.3	9.3	8.8	-2.3
Finqa	Acc	71.1*	31.4	0.7	1.7	79.7*	5.9	2.6	0	29.5	27.3	25.6	<b>36.2</b>	+4.8
<b>Average</b>		57.6*	26.1	16.5	8.4	54.2*	28.6*	3.7	0	35.3	35.5	34.0	<b>38.4</b>	+8.2

Table 2: The overall evaluation results of our model against other baselines. 7B-M was trained with Mistral-7B as the base model. Cells with “-” in the held-in part mean that the model did not train on this dataset, and results are not comparable. USKG models are overfit to the held-in dataset labels, and thus cannot generalize comparably. Cells in the held-out section with “\*” are held-in results. SoTA results are copied from the original papers for reference. ST refers to the single-task fine-tuning result of CodeLlama-Instruct-7B on each dataset. BASE and BASE-M refer to the 1-shot performance of CodeLlama-Instruct-7B, and Mistral-7B-Instruct-v0.2 respectively.  $\Delta$  refers to the difference between StructLM and the best known result. **score** denotes the state-of-the-art score on specific tasks. All StructLM held-out results are 0-shot. Specifications as to how SoTA results are selected are given in Table 4.

still surpassing USKG models on 16 of 18 datasets. It is worth noting that although the single-task (ST) models are more than double the size in parameters compared to USKG, they do not perform much better on average. This fact indicates that there may be significant unused model capacity that can be better utilized via more effective training regimes, such as our instruction tuning. Regarding FLAN-UL2-20B (Tay et al., 2023b), which was also extensively trained on structure knowledge grounding tasks, StructLM outperforms it on 7 of the 9 mutually held-in datasets. Our results on held-in datasets (Tabfact and FeTaQA) are on par with TableLlama (Zhang et al., 2023), which is an LLM pre-trained on 2.6M table understanding tasks. We are vastly outperforming TableLlama on the held-out datasets.

**Held-out Results** On held out tasks, StructLM shows strong generalization performance, outperforming ChatGPT on 5 of 6 tasks. These novel tasks contain remarkably different data format than the training dataset. For example, FinQA (Chen et al., 2021) requires the model to generate a mathematical expression to answer a question in the financial domain, while BIRD (Li et al., 2023b) requires the model to interface with multiple databases. These skills are not exactly covered in the training set, thus they put high demand for the models’ generalization capabilities. Such generalization capabilities are non-existent in USKG models as they are task-specific. Another table foundation model TableLlama (Zhang et al., 2023) would also fail on most of the held-out tasks with performance close to zero. In contrast, StructLM is much more capable of generalization on these novel tasks. The average improvement over TableLlama is 35%. Compared to another strong foundation




Figure 4: Effect of different pretraining curricula on SKG finetuning performance in relevant task groupings. We can observe the advantages of CodeLlama over the others.

Purpose	Train	Eval	FT	Result
Schema task transfer	Spider, SParC, Logic2Text	Logic2Text	89.47	89.93
KT task transfer	CompWebQ, WebQSP, GrailQa, DART	DART	60.28	60.34
Table task transfer	FetaQA, HybridQA, WikiTQ, TabMWP, ToTTo, MMQA, WikiSQL, KVRet, Tab Fact, Feverous, Infotabs	TabFact, Feverous, Infotabs	75.46	80.81
Summ. data type transfer	ToTTo, DART	DART	60.28	61.42
QA data type transfer	CompWebQ, WikiSQL	WikiSQL	85.49	86.36

Table 3: Cross task and cross datatype transfer results. FT is an average of single-task performance over the datasets in the Eval column.




Figure 5: Effect of general instruction-following data on averaged held-out SKG dataset performance. Performance is measured as the average over evaluation metrics across all tasks within held-in or held-out groups.

structure-knowledge-grounding model *Flan-UL2 20B* (Tay et al., 2023b), we also outperform it on held-out tasks by an average of 10%.

Additionally, our results in Table 2 support that the Mistral base model has stronger generalization and ability to handle SKG tasks than CodeLlama, with about a 10% advantage on metrics over all tasks. We can see that this strength transfers to the fine-tuned versions of the models.

### 5 Ablation Studies

**Effect of base model pretraining data.** We ablate our choice of base model, CodeLlama-7b-Instruct, by finetuning the unspecialized Llama2-7b base model and Llemma, which is further pretrained on mathematical texts (Azerbayev et al., 2023b). Intuitively, one might guess that coding ability has the most transferability to performance on the types of SKG tasks we are studying due to the symbolic nature of programming languages and code-



writing scenarios. However, it is possible that other types of pretraining to boost reasoning ability, such as math, have even greater transferability.

Our ablation results in Table 6 can be broken down into groupings of tasks, as in Figure 4. Models pretrained on code indeed perform slightly better, and these gains are not necessarily limited to tasks which explicitly involve a grammatically regular input, or require the generation of code. Math pretraining does seem to improve the performance of the Llama2 base model, but not by as much as code pretraining. Overall, it seems that code pretraining is a useful step in training a performant model in this SKG setting, which may be due to conceptual similarity on certain tasks.

**Effect of general instruction data in mixture** As we see in Figure 5, the held-in performance is relatively unaffected by the added general examples, but held-out performance improves significantly with more general data. Empirically, we also observe that when training a large volume of task-specific input and output formats, the model becomes less capable of following instructions on new tasks in a zero-shot setting. We hypothesize that training on this general mixture helps zero-shot performance because it can reduce overfitting to the task formats in the training set.

**Cross-task and cross-format transferability** We ablate the transferability of performance between input structure knowledge types and between output task types. To test this, we train a number of tasks together and compare them to their single-task finetuned models. Our results are indicative that there is cross-task transferability of performance occurring. On tables, we see this effect clearly. This may be explained by the volume and variety of table tasks included in the training mix. In schema and knowledge triples, the performance improvement is not pronounced, perhaps due to the limited size of those datasets. Nevertheless, evidence supports that diversifying tasks on the same structured knowledge type benefits performance.

On the other hand, we see that finetuning on different datatypes with the same task (i.e. summarization) also yields benefits to performance. On the summarization and question-answering (QA) experiments, we train on both tabular and knowledge graph data. We evaluate summarization with DART and QA with WikiSQL. We see that in both cases, the extra dataset yielded about 1% improvement. Considering that the added datasets in each case organize information in a completely different way, this result suggests diversifying data types for similar tasks do indeed benefit each other as well.

## 6 Discussion

We argue that SKG is an important capability for future language models. We have seen through our experiments on ChatGPT and the Llama2 family that there is significant room for improvement. We found that we could produce a strong model by focused instruction-tuning on SKG tasks, however, we also observe that the performance difference between 7B to 34B StructLM models was not dramatic. This raises a concern about the state of SKG data: could we be approaching a performance ceiling? Combined with the fact that we were able to outperform UL2-20b, a much larger model, with our 7B model on 3 tasks, it seems that LLMs at various scales are struggling with SKG capabilities.

Indeed, grounding to structured knowledge directly in a model’s input represents a challenge in reasoning and input sensitivity. However, it has a wide range of potential benefits. To meaningfully improve SKG capability, we propose that future work may explore continued pretraining of open foundation models on more structured data formats. Similar to current attempts at code or math pretraining, it is possible that pretraining models on text interleaved with tables or other types of regular data formatting will help us move towards establishing SKG as a foundational model capability.

## 7 Conclusion

In this paper, we explore the current capabilities of open language models on structured knowledge grounding tasks. We show that LLMs are currently weak at SKG tasks. To address this gap, we construct an instruction-tuning dataset mixture of 1.1M examples and release models that and achieve SOTA on 7 of 18 held-in tasks, and that outperform strong existing models such as TableLlama and Flan UL2 on held-out tasks. We also study the effects of various factors that influence the performance of our model on these task types. We hope that our work provides an updated understanding of what is achievable in the SKG domain, and can serve as a strong baseline for future improvements.

## References

- Rami Aly, Zhijiang Guo, Michael Sejr Schlichtkrull, James Thorne, Andreas Vlachos, Christos Christodoulopoulos, Oana Cocarascu, and Arpit Mittal. FEVEROUS: Fact extraction and VERification over unstructured and structured information. 2021.
- Zhangir Azerbayev, Hailey Schoelkopf, Keiran Paster, Marco Dos Santos, Stephen McAleer, Albert Q. Jiang, Jia Deng, Stella Biderman, and Sean Welleck. Llemma: An open language model for mathematics. *CoRR*, abs/2310.10631, 2023a.
- Zhangir Azerbayev, Hailey Schoelkopf, Keiran Paster, Marco Dos Santos, Stephen McAleer, Albert Q. Jiang, Jia Deng, Stella Biderman, and Sean Welleck. Llemma: An open language model for mathematics, 2023b.
- Junwei Bao, Duyu Tang, Nan Duan, Zhao Yan, Yuanhua Lv, Ming Zhou, and Tiejun Zhao. Table-to-text: Describing table region with natural language. *Proceedings of the AAAI Conference on Artificial Intelligence*, 32(1), April 2018. ISSN 2159-5399. doi: 10.1609/aaai.v32i1.11944. URL <http://dx.doi.org/10.1609/aaai.v32i1.11944>.
- Paweł Budzianowski, Tsung-Hsien Wen, Bo-Hsiang Tseng, Iñigo Casanueva, Stefan Ultes, Osman Ramadan, and Milica Gašić. MultiWOZ - a large-scale multi-domain Wizard-of-Oz dataset for task-oriented dialogue modelling. In Ellen Riloff, David Chiang, Julia Hockenmaier, and Jun'ichi Tsujii (eds.), *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pp. 5016–5026, Brussels, Belgium, October–November 2018. Association for Computational Linguistics. doi: 10.18653/v1/D18-1547. URL <https://aclanthology.org/D18-1547>.
- Wenhu Chen, Hongmin Wang, Jianshu Chen, Yunkai Zhang, Hong Wang, Shiyang Li, Xiyu Zhou, and William Yang Wang. Tabfact: A large-scale dataset for table-based fact verification. In *International Conference on Learning Representations*, 2019.
- Wenhu Chen, Hanwen Zha, Zhiyu Chen, Wenhan Xiong, Hong Wang, and William Wang. Hybridqa: A dataset of multi-hop question answering over tabular and textual data. *Findings of EMNLP 2020*, 2020a.
- Zhiyu Chen, Wenhu Chen, Hanwen Zha, Xiyu Zhou, Yunkai Zhang, Sairam Sundaresan, and William Yang Wang. Logic2Text: High-fidelity natural language generation from logical forms. In Trevor Cohn, Yulan He, and Yang Liu (eds.), *Findings of the Association for Computational Linguistics: EMNLP 2020*, pp. 2096–2111, Online, November 2020b. Association for Computational Linguistics. doi: 10.18653/v1/2020.findings-emnlp.190. URL <https://aclanthology.org/2020.findings-emnlp.190>.
- Zhiyu Chen, Wenhu Chen, Charese Smiley, Sameena Shah, Iana Borova, Dylan Langdon, Reema Moussa, Matt Beane, Ting-Hao Huang, Bryan Routledge, and William Yang Wang. Finqa: A dataset of numerical reasoning over financial data. *Proceedings of EMNLP 2021*, 2021.
- Zhoujun Cheng, Tianbao Xie, Peng Shi, Chengzu Li, Rahul Nadkarni, Yushi Hu, Caiming Xiong, Dragomir Radev, Mari Ostendorf, Luke Zettlemoyer, et al. Binding language models in symbolic languages. In *The Eleventh International Conference on Learning Representations*, 2022.

- Jordan Clive, Kris Cao, and Marek Rei. Control prefixes for parameter-efficient text generation. *arXiv preprint arXiv:2110.08329*, 2021.
- Yinpei Dai, Hangyu Li, Yongbin Li, Jian Sun, Fei Huang, Luo Si, and Xiaodan Zhu. Preview, attend and review: Schema-aware curriculum learning for multi-domain dialog state tracking, 2021.
- Mihail Eric, Lakshmi Krishnan, Francois Charette, and Christopher D. Manning. Key-value retrieval networks for task-oriented dialogue. In *Proceedings of the 18th Annual SIGdial Meeting on Discourse and Dialogue*. Association for Computational Linguistics, 2017. doi: 10.18653/v1/w17-5506. URL <http://dx.doi.org/10.18653/v1/W17-5506>.
- Yu Gu, Sue Kase, Michelle Vanni, Brian Sadler, Percy Liang, Xifeng Yan, and Yu Su. Beyond iid: three levels of generalization for question answering on knowledge bases. In *Proceedings of the Web Conference 2021*, pp. 3477–3488. ACM.
- Deepak Gupta, Surabhi Kumari, Asif Ekbal, and Pushpak Bhattacharyya. MMQA: A multi-domain multi-lingual question-answering framework for English and Hindi. In Nicoletta Calzolari, Khalid Choukri, Christopher Cieri, Thierry Declerck, Sara Goggi, Koiti Hasida, Hitoshi Isahara, Bente Maegaard, Joseph Mariani, H el ene Mazo, Asuncion Moreno, Jan Odijk, Stelios Piperidis, and Takenobu Tokunaga (eds.), *Proceedings of the Eleventh International Conference on Language Resources and Evaluation (LREC 2018)*, Miyazaki, Japan, May 2018. European Language Resources Association (ELRA). URL <https://aclanthology.org/L18-1440>.
- Vivek Gupta, Maitrey Mehta, Pegah Nokhiz, and Vivek Srikumar. INFOTABS: Inference on tables as semi-structured data. In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pp. 2309–2324, Online, July 2020a. Association for Computational Linguistics. URL <https://www.aclweb.org/anthology/2020.acl-main.210>.
- Vivek Gupta, Maitrey Mehta, Pegah Nokhiz, and Vivek Srikumar. Infotabs: Inference on tables as semi-structured data. *arXiv preprint arXiv:2005.06117*, 2020b.
- Dan Hendrycks, Collin Burns, Saurav Kadavath, Akul Arora, Steven Basart, Eric Tang, Dawn Song, and Jacob Steinhardt. Measuring mathematical problem solving with the math dataset, 2021.
- Mohit Iyyer, Wen-tau Yih, and Ming-Wei Chang. Search-based neural structured learning for sequential question answering. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 1821–1831, Vancouver, Canada, July 2017. Association for Computational Linguistics. doi: 10.18653/v1/P17-1167. URL <https://aclanthology.org/P17-1167>.
- Jinhao Jiang, Kun Zhou, Zican Dong, Keming Ye, Wayne Xin Zhao, and Ji-Rong Wen. StructGPT: A General Framework for Large Language Model to Reason over Structured Data, October 2023. URL <http://arxiv.org/abs/2305.09645>. arXiv:2305.09645 [cs].
- Sung-Min Lee, Eunhwan Park, Daeryong Seo, Donghyeon Jeon, Inho Kang, and Seung-Hoon Na. MAFiD: Moving average equipped fusion-in-decoder for question answering over tabular and textual data. In Andreas Vlachos and Isabelle Augenstein (eds.), *Findings of the Association for Computational Linguistics: EACL 2023*, pp. 2337–2344, Dubrovnik, Croatia, May 2023. Association for Computational Linguistics. doi: 10.18653/v1/2023.findings-eacl.177. URL <https://aclanthology.org/2023.findings-eacl.177>.
- Haoran Li, Abhinav Arora, Shuohui Chen, Anchit Gupta, Sonal Gupta, and Yashar Mehdad. MTOP: A comprehensive multilingual task-oriented semantic parsing benchmark. In Paola Merlo, Jorg Tiedemann, and Reut Tsarfay (eds.), *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pp. 2950–2962, Online, April 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.eacl-main.257. URL <https://aclanthology.org/2021.eacl-main.257>.

- Haoyang Li, Jing Zhang, Cuiping Li, and Hong Chen. Resdsq: Decoupling schema linking and skeleton parsing for text-to-sql. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 37, pp. 13067–13075, 2023a.
- Jinyang Li, Binyuan Hui, Ge Qu, Binhua Li, Jiayi Yang, Bowen Li, Bailin Wang, Bowen Qin, Ruiying Geng, Nan Huo, Xuanhe Zhou, Chenhao Ma, Guoliang Li, Kevin C. C. Chang, Fei Huang, Reynold Cheng, and Yongbin Li. Can llm already serve as a database interface? a big bench for large-scale database grounded text-to-sqls, 2023b.
- Raymond Li, Loubna Ben Allal, Yangtian Zi, Niklas Muennighoff, Denis Kocetkov, Chenghao Mou, Marc Marone, Christopher Akiki, Jia Li, Jenny Chim, Qian Liu, Evgenii Zheltonozhskii, Terry Yue Zhuo, Thomas Wang, Olivier Dehaene, Mishig Davaadorj, Joel Lamy-Poirier, João Monteiro, Oleh Shliazhko, Nicolas Gontier, Nicholas Meade, Armel Zebaze, Ming-Ho Yee, Logesh Kumar Umapathi, Jian Zhu, Benjamin Lipkin, Muh-tasham Oblokulov, Zhiruo Wang, Rudra Murthy V, Jason Stillerman, Siva Sankalp Patel, Dmitry Abulkhanov, Marco Zocca, Manan Dey, Zhihan Zhang, Nour Moustafa-Fahmy, Ur-vashi Bhattacharyya, Wenhao Yu, Swayam Singh, Sasha Luccioni, Paulo Villegas, Maxim Kunakov, Fedor Zhdanov, Manuel Romero, Tony Lee, Nadav Timor, Jennifer Ding, Claire Schlesinger, Hailey Schoelkopf, Jan Ebert, Tri Dao, Mayank Mishra, Alex Gu, Jennifer Robinson, Carolyn Jane Anderson, Brendan Dolan-Gavitt, Danish Contractor, Siva Reddy, Daniel Fried, Dzmitry Bahdanau, Yacine Jernite, Carlos Muñoz Ferrandis, Sean Hughes, Thomas Wolf, Arjun Guha, Leandro von Werra, and Harm de Vries. Starcoder: may the source be with you! *CoRR*, abs/2305.06161, 2023c.
- Shujie Li, Liang Li, Ruiying Geng, Min Yang, Binhua Li, Guanghu Yuan, Wanwei He, Shao Yuan, Can Ma, Fei Huang, and Yongbin Li. Unifying structured data as graph for data-to-text pre-training. *ArXiv*, abs/2401.01183, 2024. URL <https://api.semanticscholar.org/CorpusID:266725545>.
- Yujia Li, David H. Choi, Junyoung Chung, Nate Kushman, Julian Schrittwieser, Rémi Leblond, Tom Eccles, James Keeling, Felix Gimeno, Agustin Dal Lago, Thomas Hubert, Peter Choy, Cyprien de Masson d’Autume, Igor Babuschkin, Xinyun Chen, Po-Sen Huang, Johannes Welbl, Sven Gowal, Alexey Cherepanov, James Molloy, Daniel J. Mankowitz, Esme Sutherland Robson, Pushmeet Kohli, Nando de Freitas, Koray Kavukcuoglu, and Oriol Vinyals. Competition-level code generation with alphacode. *CoRR*, abs/2203.07814, 2022.
- Wing Lian, Guan Wang, Bley Goodson, Eugene Pentland, Austin Cook, Chanvichet Vong, and “Teknium”. Slimorca: An open dataset of gpt-4 augmented flan reasoning traces, with verification, 2023. URL <https://huggingface.co/Open-Orca/SlimOrca>.
- Guang Liu, Jie Yang, and Ledell Wu. Ptab: Using the pre-trained language model for modeling tabular data. *CoRR*, abs/2209.08060, 2022.
- Qian Liu, Bei Chen, Jiaqi Guo, Morteza Ziyadi, Zeqi Lin, Weizhu Chen, and Jian-Guang Lou. Tapex: Table pre-training via learning a neural sql executor, 2021.
- Pan Lu, Liang Qiu, Kai-Wei Chang, Ying Nian Wu, Song-Chun Zhu, Tanmay Rajpurohit, Peter Clark, and Ashwin Kalyan. Dynamic prompt learning via policy gradient for semi-structured mathematical reasoning. In *International Conference on Learning Representations (ICLR)*, 2023.
- Haipeng Luo, Qingfeng Sun, Can Xu, Pu Zhao, Jianguang Lou, Chongyang Tao, Xiubo Geng, Qingwei Lin, Shifeng Chen, and Dongmei Zhang. Wizardmath: Empowering mathematical reasoning for large language models via reinforced evol-instruct. *arXiv preprint arXiv:2308.09583*, 2023a.
- Haoran Luo, Haihong E, Zichen Tang, Shiyao Peng, Yikai Guo, Wentai Zhang, Chenghao Ma, Guanting Dong, Meina Song, and Wei Lin. Chatkbqa: A generate-then-retrieve framework for knowledge base question answering with fine-tuned large language models, 2023b.

- Ziyang Luo, Can Xu, Pu Zhao, Qingfeng Sun, Xiubo Geng, Wenxiang Hu, Chongyang Tao, Jing Ma, Qingwei Lin, and Daxin Jiang. Wizardcoder: Empowering code large language models with evol-instruct. *arXiv preprint arXiv:2306.08568*, 2023c.
- Linyong Nan, Dragomir Radev, Rui Zhang, Amrit Rau, Abhinand Sivaprasad, Chiachun Hsieh, Xiangru Tang, Aadit Vyas, Neha Verma, Pranav Krishna, Yangxiaokang Liu, Nadia Irwanto, Jessica Pan, Faiaz Rahman, Ahmad Zaidi, Mutethia Mutuma, Yasin Tarabar, Ankit Gupta, Tao Yu, Yi Chern Tan, Xi Victoria Lin, Caiming Xiong, Richard Socher, and Nazneen Fatema Rajani. DART: Open-domain structured data record to text generation. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pp. 432–447, Online, June 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.naacl-main.37. URL <https://aclanthology.org/2021.naacl-main.37>.
- Linyong Nan, Chiachun Hsieh, Ziming Mao, Xi Victoria Lin, Neha Verma, Rui Zhang, Wojciech Kryściński, Hailey Schoelkopf, Riley Kong, Xiangru Tang, Mutethia Mutuma, Ben Rosand, Isabel Trindade, Renusree Bandaru, Jacob Cunningham, Caiming Xiong, Dragomir Radev, and Dragomir Radev. Fetaqa: Free-form table question answering. *Transactions of the Association for Computational Linguistics*, 10:35–49, 2022. ISSN 2307-387X. doi: 10.1162/tacl.a.00446. URL <http://dx.doi.org/10.1162/tacl.a.00446>.
- Ankur P Parikh, Xuezhi Wang, Sebastian Gehrmann, Manaal Faruqui, Bhuwan Dhingra, Diyi Yang, and Dipanjan Das. ToTTo: A controlled table-to-text generation dataset. In *Proceedings of EMNLP*, 2020.
- Panupong Pasupat and Percy Liang. Compositional semantic parsing on semi-structured tables. In Chengqing Zong and Michael Strube (eds.), *Proceedings of the 53rd Annual Meeting of the Association for Computational Linguistics and the 7th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pp. 1470–1480, Beijing, China, July 2015. Association for Computational Linguistics. doi: 10.3115/v1/P15-1142. URL <https://aclanthology.org/P15-1142>.
- Jiexing Qi, Jingyao Tang, Ziwei He, Xiangpeng Wan, Yu Cheng, Chenghu Zhou, Xinbing Wang, Quanshi Zhang, and Zhouhan Lin. RASAT: integrating relational structures into pretrained seq2seq model for text-to-sql. In *EMNLP*, pp. 3215–3229. Association for Computational Linguistics, 2022a.
- Jiexing Qi, Jingyao Tang, Ziwei He, Xiangpeng Wan, Yu Cheng, Chenghu Zhou, Xinbing Wang, Quanshi Zhang, and Zhouhan Lin. Rasat: Integrating relational structures into pretrained seq2seq model for text-to-sql. *arXiv preprint arXiv:2205.06983*, 2022b.
- Cheng Qian, Chi Han, Yi R. Fung, Yujia Qin, Zhiyuan Liu, and Heng Ji. Creator: Tool creation for disentangling abstract and concrete reasoning of large language models, 2023.
- Yujia Qin, Shihao Liang, Yining Ye, Kunlun Zhu, Lan Yan, Yaxi Lu, Yankai Lin, Xin Cong, Xiangru Tang, Bill Qian, Sihan Zhao, Lauren Hong, Runchu Tian, Ruobing Xie, Jie Zhou, Mark Gerstein, Dahai Li, Zhiyuan Liu, and Maosong Sun. Toolllm: Facilitating large language models to master 16000+ real-world apis, 2023.
- Jeff Rasley, Samyam Rajbhandari, Olatunji Ruwase, and Yuxiong He. Deepspeed: System optimizations enable training deep learning models with over 100 billion parameters. In *Proceedings of the 26th ACM SIGKDD International Conference on Knowledge Discovery & Data Mining, KDD '20*, pp. 3505–3506, New York, NY, USA, 2020. Association for Computing Machinery. ISBN 9781450379984. doi: 10.1145/3394486.3406703. URL <https://doi.org/10.1145/3394486.3406703>.
- Baptiste Rozière, Jonas Gehring, Fabian Gloeckle, Sten Sootla, Itai Gat, Xiaoqing Ellen Tan, Yossi Adi, Jingyu Liu, Tal Remez, Jérémy Rapin, Artyom Kozhevnikov, Ivan Evtimov, Joanna Bitton, Manish Bhatt, Cristian Canton-Ferrer, Aaron Grattafiori, Wenhan Xiong, Alexandre Défossez, Jade Copet, Faisal Azhar, Hugo Touvron, Louis Martin, Nicolas Usunier, Thomas Scialom, and Gabriel Synnaeve. Code llama: Open foundation models for code. *CoRR*, abs/2308.12950, 2023.



- Chang Shu, Yusen Zhang, Xiangyu Dong, Peng Shi, Tao Yu, and Rui Zhang. Logic-consistency text generation from semantic parses. In Chengqing Zong, Fei Xia, Wenjie Li, and Roberto Navigli (eds.), *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pp. 4414–4426, Online, August 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.findings-acl.388. URL <https://aclanthology.org/2021.findings-acl.388>.
- Yiheng Shu and Zhiwei Yu. Data distribution bottlenecks in grounding language models to knowledge bases, 2023.
- Jiashuo Sun, Hang Zhang, Chen Lin, Yeyun Gong, Jian Guo, and Nan Duan. Apollo: An optimized training approach for long-form numerical reasoning. *arXiv preprint arXiv:2212.07249*, 2022.
- Alon Talmor and Jonathan Berant. The web as a knowledge-base for answering complex questions. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*. Association for Computational Linguistics, 2018. doi: 10.18653/v1/n18-1059. URL <http://dx.doi.org/10.18653/v1/N18-1059>.
- Yi Tay, Mostafa Dehghani, Vinh Q. Tran, Xavier Garcia, Jason Wei, Xuezhi Wang, Hyung Won Chung, Dara Bahri, Tal Schuster, Huaixiu Steven Zheng, Denny Zhou, Neil Houlsby, and Donald Metzler. UL2: unifying language learning paradigms. In *ICLR*. OpenReview.net, 2023a.
- Yi Tay, Mostafa Dehghani, Vinh Q. Tran, Xavier Garcia, Jason Wei, Xuezhi Wang, Hyung Won Chung, Siamak Shakeri, Dara Bahri, Tal Schuster, Huaixiu Steven Zheng, Denny Zhou, Neil Houlsby, and Donald Metzler. UL2: Unifying language learning paradigms, 2023b.
- Hugo Touvron, Louis Martin, Kevin Stone, Peter Albert, Amjad Almahairi, Yasmine Babaei, Nikolay Bashlykov, Soumya Batra, Prajjwal Bhargava, Shruti Bhosale, Dan Bikel, Lukas Blecher, Cristian Canton-Ferrer, Moya Chen, Guillem Cucurull, David Esiobu, Jude Fernandes, Jeremy Fu, Wenyin Fu, Brian Fuller, Cynthia Gao, Vedanuj Goswami, Naman Goyal, Anthony Hartshorn, Saghar Hosseini, Rui Hou, Hakan Inan, Marcin Kardas, Viktor Kerkez, Madian Khabsa, Isabel Kloumann, Artem Korenev, Punit Singh Koura, Marie-Anne Lachaux, Thibaut Lavril, Jenya Lee, Diana Liskovich, Yinghai Lu, Yuning Mao, Xavier Martinet, Todor Mihaylov, Pushkar Mishra, Igor Molybog, Yixin Nie, Andrew Poulton, Jeremy Reizenstein, Rashi Rungta, Kalyan Saladi, Alan Schelten, Ruan Silva, Eric Michael Smith, Ranjan Subramanian, Xiaoqing Ellen Tan, Binh Tang, Ross Taylor, Adina Williams, Jian Xiang Kuan, Puxin Xu, Zheng Yan, Iliyan Zarov, Yuchen Zhang, Angela Fan, Melanie Kambadur, Sharan Narang, Aurélien Rodriguez, Robert Stojnic, Sergey Edunov, and Thomas Scialom. Llama 2: Open foundation and fine-tuned chat models. *CoRR*, abs/2307.09288, 2023.
- Jason Wei, Maarten Bosma, Vincent Y. Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew M. Dai, and Quoc V. Le. Finetuned language models are zero-shot learners. In *ICLR*. OpenReview.net, 2022.
- Dongling Xiao, Linzheng Chai, Qian-Wen Zhang, Zhao Yan, Zhoujun Li, and Yunbo Cao. Cqr-sql: Conversational question reformulation enhanced context-dependent text-to-sql parsers, 2022.
- 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. Unifedskg: Unifying and multi-tasking structured knowledge grounding with text-to-text language models. *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, 2022. doi: 10.18653/v1/2022.emnlp-main.39. URL <http://dx.doi.org/10.18653/v1/2022.emnlp-main.39>.

- Can Xu, Qingfeng Sun, Kai Zheng, Xiubo Geng, Pu Zhao, Jiazhan Feng, Chongyang Tao, and Daxin Jiang. Wizardlm: Empowering large language models to follow complex instructions. *arXiv preprint arXiv:2304.12244*, 2023.
- Kuan Xu, Yongbo Wang, Yongliang Wang, Zujie Wen, and Yang Dong. Sead: End-to-end text-to-sql generation with schema-aware denoising. *arXiv preprint arXiv:2105.07911*, 2021.
- Yunhu Ye, Binyuan Hui, Min Yang, Binhua Li, Fei Huang, and Yongbin Li. Large language models are versatile decomposers: Decomposing evidence and questions for table-based reasoning. In *Proceedings of the 46th International ACM SIGIR Conference on Research and Development in Information Retrieval*, pp. 174–184, 2023.
- Tao Yu, Rui Zhang, Kai Yang, Michihiro Yasunaga, Dongxu Wang, Zifan Li, James Ma, Irene Li, Qingning Yao, Shanelle Roman, Zilin Zhang, and Dragomir Radev. Spider: A large-scale human-labeled dataset for complex and cross-domain semantic parsing and text-to-sql task. In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*. Association for Computational Linguistics, 2018. doi: 10.18653/v1/d18-1425. URL <http://dx.doi.org/10.18653/v1/D18-1425>.
- Tao Yu, Rui Zhang, Michihiro Yasunaga, Yi Chern Tan, Xi Victoria Lin, Suyi Li, Heyang Er, Irene Li, Bo Pang, Tao Chen, Emily Ji, Shreya Dixit, David Proctor, Sungrok Shim, Jonathan Kraft, Vincent Zhang, Caiming Xiong, Richard Socher, and Dragomir Radev. Sparc: Cross-domain semantic parsing in context. In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*. Association for Computational Linguistics, 2019. doi: 10.18653/v1/p19-1443. URL <http://dx.doi.org/10.18653/v1/P19-1443>.
- Tao Yu, Chien-Sheng Wu, Xi Victoria Lin, Yi Chern Tan, Xinyi Yang, Dragomir Radev, Caiming Xiong, et al. Grappa: Grammar-augmented pre-training for table semantic parsing. In *International Conference on Learning Representations*, 2020.
- Tianshu Zhang, Xiang Yue, Yifei Li, and Huan Sun. TableLlama: Towards Open Large Generalist Models for Tables, November 2023. URL <http://arxiv.org/abs/2311.09206>. arXiv:2311.09206 [cs].
- Yilun Zhao, Yunxiang Li, Chenying Li, and Rui Zhang. MultihierTT: Numerical reasoning over multi hierarchical tabular and textual data. In *ACL (1)*, pp. 6588–6600. Association for Computational Linguistics, 2022.
- Victor Zhong, Caiming Xiong, and Richard Socher. Seq2sql: Generating structured queries from natural language using reinforcement learning. *CoRR*, abs/1709.00103, 2017.

## A SoTA Results

Dataset	Metric	Split	SOTA Score	Best Performing Model
TabMWP	Acc	test	94.7	CREATOR (ChatGPT)(Qian et al., 2023)
ToTTo	BLEU	test	49.9	UniK2G (Li et al., 2024)
GrailQA	EM	val	77.1	TIARA + GAIN (T5-3B) (Shu & Yu, 2023)
SQL2Text	Blec	test	94.78	UnifiedSKG
MMQA	F1	test	85.28	UnifiedSKG
Spider	EM	dev	80.5	RESDSL (Li et al., 2023a)
KVRet	All Micro	test	67.88	UnifiedSKG
HybridQA	Acc	dev	68.4	S3HQA(Lee et al., 2023)
SParC	EM	test	68.2	CQR-SQL(Qi et al., 2022b)
CompWebQ	Acc	test	76.8	ChatKBQALuo et al. (2023b)
TabFact	Acc	test_small	93.0	Dater(Ye et al., 2023)
WikiTQ	All Ex	test	65.9	Dater(Ye et al., 2023)
WikiSQL	All Ex	test	93.0	SeaD+Execution-Guided Decoding(Xu et al., 2021)
FeTaQA	BLEU	test	39.0	TableLlama (Zhang et al., 2023)
Feverous	Acc	dev	85.6	FLAN UL2 20b(Tay et al., 2023b)
MultiWOZ	Joint Acc	test	60.6	TripPy + SaCLog(Dai et al., 2021)
Dart	BLEU	test	52.0	Control Prefixes (T5-large)(Clive et al., 2021)
Logic2Text	Blec	test	95.3	UnifiedSKG
MTOP	EM	test	87.5	FLAN UL2 20b(Tay et al., 2023b)
BIRD	Acc	dev	36.6	ChatGPT + COT(Li et al., 2023b)
CoSQL	EM	test	58.3	CQR-SQL(Xiao et al., 2022)
SQA	Acc	test	70.5	FLAN UL2 20b(Tay et al., 2023b)
Infotabs	Acc	dev	75.6	Infotabs paper(Gupta et al., 2020b)
WikiTableText	BLEU	test	33.7	UnifiedK2G(Li et al., 2024)
Finqa	Acc	private_test	71.1	APOLLO(Sun et al., 2022)

Table 4: Specification of SoTA scores and their sources on the most prevalent metrics used for assessment.

Across the datasets that we evaluate, the SoTA scores are chosen based on methods that do not use agent based methods on models as large as GPT-4 for a fairer comparison. Across these datasets, we can see that many of the best performing methods are purpose-built for the type of data structure used in the task. For example, RESDSL focuses exclusively on controlled SQL generation, DATER uses SQL-based reasoning for tabular tasks, and S3HQA focuses on table and text multi-hop QA. Compared to these methods, the performance of StructLM can be seen a strong baseline to determine if these domain-specific designs yield real benefits.

## B SlimOrca Dataset Mixture Details

	Metric	0%	10%	20%	50%	57%
Held-In Datasets						
TabMWP	Acc	71.14	70.35	70.52	69.01	69.36
ToTTo	BLEU	49.78	49.51	49.47	49.31	49.38
GrailQA	EM	81.09	80.46	80.29	80.89	80.38
SQL2Text	Blec	95.07	94.39	94.49	94.97	93.81
MMQA	F1	84.26	84.31	84.11	83.40	85.15
Spider	EM	72.92	71.57	73.40	72.73	72.44
KVRet	All Micro	71.60	73.90	70.34	72.25	72.61
HybridQA	Acc	59.23	59.09	59.09	59.03	59.17
SParC	EM	63.09	62.34	63.26	64.59	61.93
CompWebQ	Acc	80.61	79.15	78.76	78.73	78.34
TabFact	Acc	83.41	81.09	81.42	80.92	80.77
WikiTQ	All Ex	50.02	48.50	49.24	48.30	50.09
WikiSQL	All Ex	87.33	86.45	86.73	86.68	88.67
FeTaQA	BLEU	36.58	37.26	36.55	36.72	36.03
Feverous	Acc	85.02	84.13	84.11	83.73	84.41
MultiWOZ	Joint Acc	54.66	54.10	53.73	53.92	54.49
Dart	BLEU	61.38	61.89	61.08	62.24	62.24
Logic2Text	Blec	88.83	89.47	89.19	90.57	88.92
MTOP	EM	82.44	81.71	81.19	80.92	81.21
Held-Out Datasets						
BIRD	Acc	21.30	22.30	22.30	23.00	22.30
CoSQL	EM	51.24	49.95	50.84	50.74	49.75
SQA	Acc	49.02	46.03	43.11	48.39	49.72
Infotabs	Acc	38.00	56.26	57.87	62.35	62.46
WikiTableText	BLEU	14.78	13.51	6.66	7.27	8.27
Finqa	Acc	19.70	24.32	27.55	25.37	27.29

Table 5: Ablation results for the mixtures of general data in the training set.

In total, we train 5 models, where the percentage represents the percent of the training data that is general. In the held out data, we see noticeable gains in generalization performance for FinQA and InfoTabs datasets. Notably, FinQA requires the generation of a python-executable math expression and InfoTabs requires an exact match to 3 previously unseen (boolean) options. WikiTableText performance seems to suffer, but is evaluated based on the BLEU score with only one target sentence. As a result, we place more emphasis on the model’s 0-shot adaptation ability to new output specifications unseen in the training data.

### C Per-dataset Pretraining Data Comparison

Tasks	Metric	Code-LM	LLaMA	Math-LM
Held-In Datasets				
TabMWP	Acc	71.14	62.96	66.5
ToTTo	BLEU	49.78	48.26	47.4
GrailQA	EM	81.09	75.72	77.66
SQL2Text	Blec	95.07	94.49	94.58
MMQA	F1	84.26	83.96	82.13
Spider	EM	72.92	65.96	71.95
KVRet	All Micro	71.6	70.36	70.03
HybridQA	Acc	59.23	59.26	57.04
SParC	EM	63.09	56.94	60.35
CompWebQ	Acc	80.61	77.31	76.6
TabFact	Acc	83.41	80.46	79.47
WikiTQ	All Ex	50.02	45.6	46.89
WikiSQL	All Ex	87.33	83.93	85.49
FeTaQA	BLEU	36.58	34.37	34.1
Feverous	Acc	85.02	83.2	82.52
MultiWOZ	Joint Acc	54.66	55.43	53.79
Dart	BLEU	61.38	61.52	61.24
Logic2Text	Blec	88.83	88.0	90.38
MTOP	EM	82.44	77.18	75.56
Held-Out Datasets				
BIRD	Acc	21.3	15.9	18.8
CoSQL	EM	51.24	42.8	48.76
SQA	Acc	49.02	37.03	49.05
Infotabs	Acc	38.0	4.44	32.54
WikiTableText	BLEU	14.78	13.0	14.82
Finqa	Acc	19.7	6.63	21.53

Table 6: Fine-grained evaluation results comparing finetuning done on different base models. Code refers to CodeLlama-Instruct-7B. Math refers to Llemma-7b. LLaMA refers to Llama2-7b.

In these fine-grained results we can compare the performance of Llemma, Llama, and CodeLlama in more detail. Llemma does seem to hold an advantage on tasks that involve math (TabMWP, FinQA). Math pretraining does not seem to benefit tabular tasks overall (WikiSQL, WikiTQ, HybridQA, MMQA, etc.), but does show advantages on SQL coding tasks (Spider, SparC). CodeLlama, however, seems to show a performance improvement over the base Llama model on not just coding tasks, but also math and tabular tasks as well. These findings may underscore the need to understand what constitutes “reasoning” in a language model.



## D Prompt Format

```
[INST] <<SYS>>
You are an AI assistant that specializes in analyzing and reasoning
over structured information. You will be given a task, optionally
with some structured knowledge input. Your answer must strictly
adhere to the output format, if specified.
<</SYS>>

{instruction} {input} [/INST]
```

Figure 6: Prompt format for all SKG examples. This formatting convention follows Llama2 [Touvron et al. \(2023\)](#). The input contains the linearized structured data, together with any other context, question or statement. The instruction specifies the task.

## E Held-Out Generation Examples

For illustration purposes, we provide examples of successful and unsuccessful responses of StructLM-13B on the FinQA held-out dataset.

### E.1 Successful StructLM-13B Response on FinQA

Use the financial data in the following passage and table to calculate the answer the question. table:

performance graph the following graph compares the yearly change in the cumulative total stockholder return for our last five full fiscal years , based upon the market price of our common stock , with the cumulative total return on a nasdaq composite index ( u.s . companies ) and a peer group , the nasdaq medical equipment-sic code 3840-3849 index , which is comprised of medical equipment companies , for that period . the performance graph assumes the investment of \$ 100 on march 31 , 2007 in our common stock , the nasdaq composite index ( u.s . companies ) and the peer group index , and the reinvestment of any and all dividends. .

col :	3/31/2007	3/31/2008	3/31/2009	3/31/2010	3/31/2011	3/31/2012
row 1 : abiomed inc	100	96.19	35.87	75.55	106.37	162.45
row 2 : nasdaq composite index	100	94.11	63.12	99.02	114.84	127.66
row 3 : nasdaq medical equipment sic code 3840-	100	82.91	41.56	77.93	94.54	74.40

this graph is not 201csoliciting material 201d under regulation 14a or 14c of the rules promulgated under the securities exchange act of 1934 , is not deemed filed with the securities and exchange commission and is not to be incorporated by reference in any of our filings under the securities act of 1933 , as amended , or the exchange act whether made before or after the date hereof and irrespective of any general incorporation language in any such filing . transfer agent american stock transfer & trust company , 59 maiden lane , new york , ny 10038 , is our stock transfer agent. .

question:

what is the roi of an investment in nasdaq composite index from march 2007 to march 2010? write down a short python expression.

---

$(99.02 - 100) / 100$

In the above example, StructLM-13B identified the relevant quantity within the given linearized table, then determined how to calculate the required ROI quantity in a python-executable expression.

## E.2 Unsuccessful StructLM-13B Response on FinQA

Use the financial data in the following passage and table to calculate the answer the question. table:

performance graph the following graph and table compares the cumulative five-year total return provided to shareholders on our class a common stock relative to the cumulative total returns of the s&p 500 index and our customized peer group . the peer group includes cboe holdings , inc. , intercontinentalexchange group , inc . and nasdaq , inc . an investment of \$ 100 ( with reinvestment of all dividends ) is assumed to have been made in our class a common stock , in the peer group and the s&p 500 index on december 31 , 2012 , and its relative performance is tracked through december 31 , 2017 . comparison of 5 year cumulative total return\* among cme group inc . , the s&p 500 index , and a peer group 12/12 12/13 12/14 12/15 12/16 cme group inc . s&p 500 peer group \* \$ 100 invested on 12/31/12 in stock or index , including reinvestment of dividends . fiscal year ending december 31 . copyright a9 2018 standard & poor 2019s , a division of s&p global . all rights reserved . the stock price performance included in this graph is not necessarily indicative of future stock price performance. .

col :	2013	2014	2015	2016	2017
row 1 : cme group inc .	\$ 164.01	\$ 194.06	\$ 208.95	\$ 279.85	\$ 370.32
row 2 : s&p 500	132.39	150.51	152.59	170.84	208.14
row 3 : peer group	176.61	187.48	219.99	249.31	323.23

unregistered sales of equity securities during the past three years there have not been any unregistered sales by the company of equity securities. .

question:

in 2017 what was the ratio of the the cme group inc . stock performamce to the s&p write down a short python expression.

---

(cme\_group\_inc\_stock\_performance / s&p\_500\_stock\_performance)

In the above example, StructLM-13B understood how to calculate the required quantity, but could not select the relevant numbers from the given table.

## Limitations

The collection process used to construct the training data for StructLM tries to include a wide variety of data types. As we have seen, there is evidence that this diversity is capable of affording transferable benefits to each dataset in the mixture. However, the tasks that we train and evaluate on are still academic datasets which have each been curated and designed for a specific purpose. While these are indeed diverse, the SKG domain relies on specific formatting and prompting conventions, which may result in our models having unnecessary specificity towards the conventions within our train set. To develop a clearer picture of how SKG performs as its own domain, we may require larger scale datasets with more heterogeneous formatting conventions. Further opportunities for training more robust SKG models may lie in increasing the diversity of structured data types in this way.

Additionally, while we have tried to evaluate our models to the best of our ability, many of the tasks of our held-out datasets measure accuracy through a heuristic matching step of a model's output. In zero or few-shot settings, it is quite challenging to exactly control the generations of an autoregressive transformer to be adherent to a certain rule or grammar, and this has been a subject of study in many of the other works cited in [Table 4](#). We note that because of this reality, poor results in zero or few-shot context may betray the existence of useful representations that the model has already learned. Without further prompting or finetuning efforts, it may be difficult to bring these capabilities to light. As such, another opportunity for improvement upon our methods may involve more flexible constrained methods of language model evaluation.