

Addressing Resource and Privacy Constraints in Semantic Parsing Through Data Augmentation

Kevin Yang^{1,*} Olivia Deng² Charles Chen²
Richard Shin² Subhro Roy² Benjamin Van Durme²
¹UC Berkeley, ²Microsoft Semantic Machines
yangk@berkeley.edu, sminfo@microsoft.com

Abstract

We introduce a novel setup for low-resource task-oriented semantic parsing which incorporates several constraints that may arise in real-world scenarios: (1) lack of similar datasets/models from a related domain, (2) inability to sample useful logical forms directly from a grammar, and (3) privacy requirements for unlabeled natural utterances. Our goal is to improve a low-resource semantic parser using utterances collected through user interactions. In this highly challenging but realistic setting, we investigate data augmentation approaches involving generating a set of structured canonical utterances corresponding to logical forms, before simulating corresponding natural language and filtering the resulting pairs. We find that such approaches are effective despite our restrictive setup: in a low-resource setting on the complex SMCaFlow calendaring dataset (Andreas et al., 2020), we observe 33% relative improvement over a non-data-augmented baseline in top-1 match.

1 Introduction

We aim to improve the performance of a semantic parser based on previous user interactions, but without making use of their direct utterances, nor any associated personal identifiable information (PII). Such privacy requirements are common in practical deployment (Kannan et al., 2016), and semantic parsers are commonly used in real-world systems such as Siri and Alexa, converting natural language into structured queries to be executed downstream (Kamath and Das, 2018).

Constructing semantic parsers can be expensive: annotating examples consisting of natural language-logical form pairs often requires trained experts. Two complementary lines of work has pursued this concern. First, several works (Zhong et al., 2020; Cao et al., 2020) tackle *low-resource* semantic parsing via approaches such as data augmentation. A

¹*: Work done during internship at Semantic Machines.

Natural	<i>When is Allison's birthday?</i>
Logical	(Yield :output (:start (singleton (:results (FindEventWrapperWithDefaults :constraint (Constraint[Event] :subject (? = #(String "Allison's birthday"))))))))
Canonical	start time of find event called something like "Allison's birthday"

Table 1: An example of natural language, logical form, and canonical form in the SMCaFlow domain. The event title, "*Allison's birthday*," is private information.

second line of work (Wang et al., 2015; Xiao et al., 2016) explores *canonical utterances*: structured language which maps directly to logical forms, but resembles natural language (Table 1). The use of canonical forms as the target of semantic parsing has shown benefits in accuracy (Shin et al., 2021; Wu et al., 2021).

We consider low-resource semantic parsing with further resource and privacy constraints which may arise in practical deployment: beyond a small gold dataset of labeled pairs, we assume only unlabeled natural utterances which must be masked for PII. Unlike many prior works, we assume that (1) we do not have a large dataset of related logical forms in a different domain, (2) we cannot sample arbitrarily many useful logical forms, and (3) we must preserve privacy of user utterances.

We propose several approaches which are compatible with our imposed restrictions, broadly following three steps: (1) generate a set of privacy-preserving canonical utterances; (2) simulate corresponding natural utterances; and (3) filter the resulting canonical-natural utterance pairs to yield additional "silver" data for training. We more than double the performance of a non-data-augmented baseline on the ATIS domain (Hemphill et al., 1990), and achieve a 33% relative improvement on the more realistic SMCaFlow domain (Andreas et al., 2020). We hope these experiments help motivate further research interest in parser improvement for realistic scenarios.

2 Semantic Parsing in Practice

Our setup assumes access exclusively to:

1. a small “seed” dataset \mathcal{D} of natural utterance with corresponding parses, and
2. a larger set \mathcal{U} of unlabeled natural utterances, for which PII must be masked before use.

In a real-world setting, one might hand-annotate the seed dataset \mathcal{D} to train a system for initial deployment, while then leveraging \mathcal{U} to refine a future version of the system.

While our setting is highly restrictive, we argue that it reflects practical constraints. For example, in practice, the grammar for logical forms—as well as the synchronous context-free grammar (SCFG) that maps them to canonical utterances—will often be written from scratch, precluding transfer learning methods which leverage a large quantity of similar data in another domain. Moreover, in complex domains, one cannot expect to *sample* useful logical forms directly from a grammar if the grammar is designed for *coverage* as in e.g., SMCaFlow (Andreas et al., 2020). Therefore, other than \mathcal{D} , the only additional data (excluding additional manual annotation) are subsequent user inputs in the form of \mathcal{U} , with PII masked to preserve privacy.

3 Related Work

Compared to prior work in low-resource semantic parsing, our task setup’s constraints require different approaches.

First, we consider semantic parsing on an entirely *new* grammar for logical forms, rather than adapting to new domains starting from a *preexisting* grammar (Zhao et al., 2019; Zhong et al., 2020; Burnyshev et al., 2021; Kim et al., 2021; Tseng et al., 2021). For example, Zhong et al. (2020) takes a natural-language-to-SQL model for one database to propose language-SQL training examples for another database.

Second, we assume one cannot sample useful canonical utterances directly from the grammar, unlike Zhong et al. (2020) and Cao et al. (2020). For example, Cao et al. (2020) use a backtranslation-esque approach leveraging large numbers of unlabeled natural and canonical utterances.

Moreover, we do not even assume direct access to unlabeled natural utterances, due to real-world privacy considerations (Kannan et al., 2016; Campagna et al., 2017). Many works on low-resource

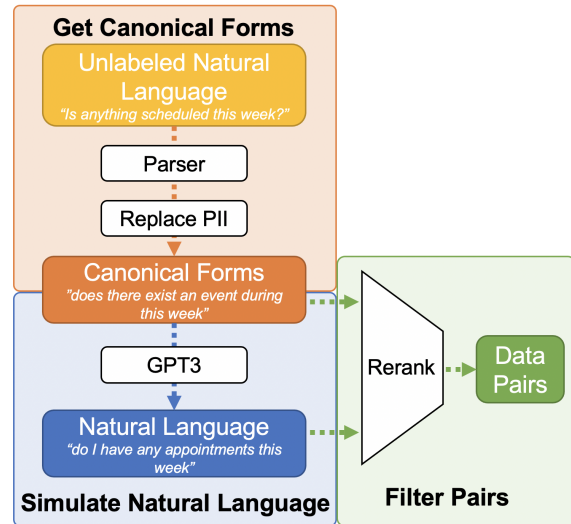


Figure 1: Illustration of one of our proposed methods for data augmentation (USER-RANK) in low-resource semantic parsing. We first obtain canonical forms from unlabeled user data using a parser trained on seed data, replacing PII. Next, we simulate corresponding natural language for the generated canonical forms. Finally, we filter the canonical-natural pairs to obtain our final silver data pairs for augmentation.

semantic parsing, such as those mentioned previously, do not consider the privacy aspect.

Nevertheless, recent work (Shin et al., 2021; Wu et al., 2021; Yin et al., 2021; Schucher et al., 2021) has demonstrated decent performance given just a small seed dataset \mathcal{D} , by combining pretrained language models with constrained decoding. For example, Shin et al. (2021) use only 300 labeled examples in the complex SMCaFlow dialogue domain (Andreas et al., 2020). However, using pretrained models to directly *generate* silver training data, with a method such as DINO (Schick and Schütze, 2021), is unsuitable in semantic parsing: the models are unaware of either the underlying grammar or the space of parsable queries. One of our contributions is to explore more effective uses of pretrained models for data augmentation in a practical semantic parsing scenario.

Finally, the detection of PII in user data is an applied topic of interest (Pilán et al., 2022), such as for safely summarizing call transcripts (Transcribe) or the automatic detection of doxing (Karimi et al., 2022). In our work we implement a solution meant as a proof of concept for our exploration, based on detecting and replacing named entities.

4 Practical Augmentation

While finetuning a pretrained model on the seed dataset \mathcal{D} can yield a reasonable parser P (Shin et al., 2021; Wu et al., 2021), we aim to increase performance via data augmentation. However, our realistic setup precludes many prior approaches. We propose to generate silver data via three main steps, shown in Figure 1: (1) generate a set \mathcal{C} of canonical utterances c , (2) simulate a set \mathcal{N} of corresponding natural utterances n , and (3) filter the resulting (c, n) pairs. We suggest multiple approaches for these steps, and benchmark their efficacy in Sec. 5. The entire procedure can be iterated multiple times as the parser improves.

4.1 Generating Canonical Utterances

First, we generate canonical utterances c . In principle, one could sample directly from a task-specific grammar, but the results may not be useful in practice (Sec. 5). The remaining options are to generate c conditioned on either unlabeled natural utterances \mathcal{U} or the seed data \mathcal{D} .

Generation conditioned on \mathcal{U} (USER). We need to mask all PII, but this is difficult to guarantee in the original natural language domain. Therefore, we first train a parser P on \mathcal{D} , and parse each utterance in \mathcal{U} to obtain a set of canonical utterances \mathcal{C}' . In the more structured domain of \mathcal{C}' we can guarantee masking and replacing all PII to yield the final set \mathcal{C} . Critically, it is not necessary that the initial \mathcal{C}' are correct parses of \mathcal{U} ; we only need a realistic distribution over canonical utterances, and the initial \mathcal{U} is no longer parallel to the final \mathcal{C} anyway due to replacing PII. Hence it is acceptable if the parser P 's errors are numerous but unbiased. In any case, the final \mathcal{C} will be somewhat tied to the true distribution of user utterances in \mathcal{U} .

Generation conditioned on \mathcal{D} (GPT). A second method of generating \mathcal{C} is SCFG-constrained decoding on an autoregressive language model,¹ prompting with the seed data \mathcal{D} . Specifically, we prompt with a random concatenation of plans from \mathcal{D} , separated by newlines. The SCFG that defines canonical utterances constrains the decoding, forcing the model to output a valid canonical utterance.

¹Ideally we would use GPT3 (Brown et al., 2020), and we do so in the ATIS domain, but API limitations in GPT3 together with the requirements of our constrained decoding force us to use GPT2-XL (Radford et al., 2019) in SMCaFlow.

4.2 Simulated Natural Utterances

For each canonical c in \mathcal{C} , we now re-generate a natural utterance n . While other methods (e.g., finetuning) are possible, here we employ a prompting approach using GPT3 (Brown et al., 2020). We use a prompt containing \mathcal{D} 's canonical-natural pairs, ending with the canonical utterance c for which we want to sample a corresponding n .

4.3 Filtering Silver Data

Many (c, n) pairs we generate may be low-quality, depending on the task and seed data \mathcal{D} available. To obtain more high-quality pairs, we simulate 20 natural utterances n for each c . We must then filter the resulting pairs, which we do based on either reranking or cycle consistency.

Reranking (RANK). We accept the best of 20 simulated n for each c , and add this (c, n) to our training data. The reranker combines two scores: (1) the log-probability that the original \mathcal{D} -trained parser P parses n back to the original canonical c , and (2) the edit distance between n and c (capped based on the length of c), which should intuitively be *maximized* to encourage linguistic diversity in the augmented data, perhaps at a small accuracy cost.

Cycle consistency (CYC). We accept a (c, n) pair if the original parser P parses n back to c . This assures the resulting pairs' quality, but may skew the distribution toward easier examples, which are less helpful in downstream training.

5 Experiments

Tasks. We evaluate on two domains, both English:

1. ATIS (Hemphill et al., 1990), a flight booking dataset. We use the Break (Wolfson et al., 2020) subset.²
2. SMCaFlow (Andreas et al., 2020), a calendaring dataset, which we view as the most complex and realistic.

In each domain, we assume a seed data \mathcal{D} of just 30 pairs, conducting several trials with different random samples of seed data to mitigate noise from this selection. We sample 300 unlabeled natural utterances \mathcal{U} from the dataset, which must be parsed to canonical forms (using the grammar and SCFG of Shin et al. (2021)) and then PII-masked

²We also ran preliminary experiments on the DROP (Dua et al., 2019) and NLVR2 (Suh et al., 2018) subsets of Break, but found that the canonical utterances were too unnatural for any method to perform reasonably (Appendix C).

before use. Our implementation of PII masking is based on recognizing and replacing named entities; see Appendix A for further details.

Methods. We evaluate several methods on each task, listed below.

1. BASE, a supervised baseline which finetunes BART (Lewis et al., 2019) on the seed \mathcal{D} following Shin et al. (2021), discarding \mathcal{U} .
2. USER-RANK, a data augmentation approach following the USER and RANK methods described in Sec. 4.1 and 4.3 respectively, and depicted in Figure 1.
3. GPT-RANK, a similar approach which generates c following GPT from Sec. 4.1 instead.
4. USER-CYC, a version which filters (c, n) pairs via cycle consistency (Sec. 4.3).
5. GRAM-RANK, a weak baseline that samples initial c directly from the grammar, which we run only on SMCaFlow since our ATIS grammar is too loosely specified for sampling.

Results. We observe that our best data augmentation methods (USER-RANK, GPT-RANK) double the performance of the baseline finetuning method BASE on ATIS, and outperform it on SMCaFlow by up to 20% relative gain (Table 2).³ Nonetheless, absolute performance remains low due to the tiny amount of seed data, although we note that the exact match metric may be unnecessarily harsh, penalizing some semantically equivalent parses.

Of interest is that GPT-RANK outperforms BASE despite using only the seed \mathcal{D} , and not extra unlabeled \mathcal{U} . Moreover, iterating the data augmentation procedure (USER-RANK-3X, GPT-RANK-3X) can further improve performance compared to BASE (relative 150% on ATIS, 33% relative on SMCaFlow), by improving the initial parser P used for parsing unlabeled \mathcal{U} or for filtering pairs (c, n) , although we observed in preliminary experiments that further iterations yielded diminishing benefits.

In contrast, USER-CYC performs poorly on ATIS, indicating that the CYC filtering is perhaps too restrictive for certain domains. Even on SMCaFlow where performance is decent in comparison, the successful cycles are overwhelmingly for relatively trivial canonical utterances (e.g., “Hello! How are you?”). We additionally observe that nearly one-third of cycles are successful, much

³Although the standard deviations appear large, the variation between trials is largely due to randomness in selecting the seed data \mathcal{D} . For example, USER-RANK is better than BASE on SMCaFlow with $p = .0004$ on a paired t -test.

Method	ATIS	SMCaFlow
BASE	6.8 ± 3.5	13.2 ± 3.4
USER-RANK	13.4 ± 4.1	15.5 ± 3.7
GPT-RANK	13.7 ± 3.2	15.9 ± 2.7
USER-CYC	6.0 ± 2.3	15.0 ± 4.0
GRAM-RANK		13.4 ± 2.8
USER-RANK-3X	17.3 ± 1.3	17.6 ± 4.6
GPT-RANK-3X	16.7 ± 3.5	16.1 ± 3.0

Table 2: *Main results on ATIS and SMCaFlow for different methods.* Top-1 parsing match percentage evaluated over 5 (ATIS) or 10 (SMCaFlow) trials on different seed datasets \mathcal{D} . For the two highest-performing methods, USER-RANK and GPT-RANK, we iterate data augmentation 3 times on SMCaFlow, yielding USER-RANK-3X and GPT-RANK-3X. USER-RANK-3X performs best overall.

more than the actual validation set accuracy of 15%, indicating that our auto-generated user utterances remain less challenging and diverse compared to real user utterances. Meanwhile, GRAM-RANK is no better than the unaugmented baseline BASE: sampling plans directly from a grammar is ineffective in a complex, realistic domain like SMCaFlow.

5.1 Analysis

We conduct additional analyses on SMCaFlow.

Reranking. First, we run ablations on reranking in USER-RANK (Table 3). While our edit distance heuristic described in Sec. 4.3 makes little difference, reranking of some form is crucial. Meanwhile, there are many possibilities for other reranking procedures.

Method	SMCaFlow
BASE	13.2 ± 3.4
USER-RANK-3X	17.6 ± 4.6
USER-NOEDITRANK-3X	17.3 ± 4.7
USER-NORANK	12.8 ± 3.5

Table 3: *SMCaFlow reranking ablations.* Since the version without reranking (USER-NORANK) is no better than the baseline, we do not iterate the data augmentation procedure. The edit distance heuristic makes little difference in this case (USER-NOEDITRANK-3X vs. USER-RANK-3X), but reranking is crucial.

Effect of Masking PII. We rerun our full pipeline for USER-RANK on SMCaFlow, removing only

the step where we resampled PII, in order to isolate the effect of doing so (Table 4, USER-RANK-KEEPPII). As one might expect, replacing PII hurts performance, albeit slightly. Of course, if PII is not a concern, then many other data augmentation schemes from prior work become possible again.

Method	SMCalFlow
USER-RANK	15.5 ± 3.7
USER-RANK-KEEPPII	16.2 ± 2.8

Table 4: *SMCalFlow* ablation where we do not resample PII. As expected, performance is slightly better if we do not need to resample PII.

Additional Seed Data. We explore using a larger seed dataset \mathcal{D} on both ATIS and SMCaFlow (90 and 100 data points respectively, instead of 30). On SMCaFlow, we observe that USER-RANK’s gains over the baseline largely disappear (Table 5). Thus, improved data augmentation methods which still yield gains with larger seed datasets are an important direction for future exploration.

Method	ATIS (90 seed)	SMCalFlow (100 seed)
BASE	21.4 ± 1.8	31.6 ± 0.3
USER-RANK	21.4 ± 1.7	31.7 ± 1.0

Table 5: *Results with more seed data.* We use a seed dataset \mathcal{D} of size 90 (ATIS) or 100 (SMCalFlow) rather than 30, with 3 trials per method. The gains from data augmentation largely disappear at this scale, so we do not do additional augmentation iterations.

Examples and Error Analysis. Finally, in Table 6 we show several SMCaFlow example parses by the baseline BASE compared to our highest-performing method USER-RANK-3X. Compared to BASE, USER-RANK-3X is often better at segmenting names (Example 1), and is also more likely to be semantically similar to the gold parse in cases where BASE is wildly incorrect (Example 2). Nevertheless, in the latter example, USER-RANK-3X is still marked wrong, suggesting that our exact match metric may somewhat underrepresent the performance of all models. Finally, both methods struggle on more complex and/or composite intents (Example 3). Additional examples illustrating these phenomena are shown in Appendix B.

Example 1	
Previous Agent	
Natural	<i>please make a meet with my doctor sarah</i>
BASE	create event with " doctor"
USER-RANK-3X	create event with " doctor sarah"
Gold Canonical	create event with " doctor sarah"
Example 2	
Previous Agent	
Natural	<i>Let me know if there’s anything else I can help you with.</i>
BASE	does there exist an event tomorrow 9 military
USER-RANK-3X	Looks good!
Gold Canonical	Thanks for your help!
Example 3	
Previous Agent	
Natural	<i>Please accept the bowling fundraiser and tell sammy I will bring refreshments.</i>
BASE	create event called " bowling fundraiser" starting month 4 11 2019 5 PM
USER-RANK-3X	create event called " bowling fundraiser"
Gold Canonical	respond Accepted with comment " I will bring refreshments" to find event called something like " bowling fundraiser"

Table 6: Example parses by the baseline BASE and our best method USER-RANK-3X on SMCaFlow. Each example contains the previous agent utterance (if it exists) and user utterance in the first two lines, followed by the BASE parse, USER-RANK-3X parse, and gold parse.

6 Conclusion

We have discussed a challenging setting for low resource semantic parsing based on real-world resource and privacy constraints. In addition to a seed dataset, the only resources allowed are unlabeled natural utterances which must be PII-masked. We observe that data augmentation approaches leveraging pretrained language models can still improve over supervised baselines which use only the seed dataset. At the same time, substantial room remains for improvement: there are many alternatives to our reranking procedure for silver data, and our method loses some effectiveness when more labeled data is provided. We hope that our exploratory observations help lay a foundation for further work in realistic data augmentation approaches for semantic parsing.

Ethical Considerations

We believe our work makes a positive impact by focusing heavily on the need for privacy considerations when exploring low-resource settings for semantic parsing. However, as our methods rely heavily on large pretrained language models such as GPT3, we may inherit similar biases which such models are known for (Brown et al., 2020).

Acknowledgements

We thank the rest of the team at Semantic Machines, as well as our anonymous reviewers, for their support and helpful comments which aided us greatly in improving the paper. We also thank the NSF for their support through a fellowship to the first author.

References

- Jacob Andreas, John Bufe, David Burkett, Charles Chen, Josh Clausman, Jean Crawford, Kate Crim, Jordan DeLoach, Leah Dorner, Jason Eisner, et al. 2020. Task-oriented dialogue as dataflow synthesis. *Transactions of the Association for Computational Linguistics*, 8:556–571.
- Tom B Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. 2020. Language models are few-shot learners. *arXiv preprint arXiv:2005.14165*.
- Pavel Burnyshev, Valentin Malykh, Andrey Bout, Ekaterina Artemova, and Irina Piontkovskaya. 2021. A single example can improve zero-shot data generation. *arXiv preprint arXiv:2108.06991*.
- Giovanni Campagna, Rakesh Ramesh, Silei Xu, Michael Fischer, and Monica S Lam. 2017. Almond: The architecture of an open, crowdsourced, privacy-preserving, programmable virtual assistant. In *Proceedings of the 26th International Conference on World Wide Web*, pages 341–350.
- Ruisheng Cao, Su Zhu, Chenyu Yang, Chen Liu, Rao Ma, Yanbin Zhao, Lu Chen, and Kai Yu. 2020. Unsupervised dual paraphrasing for two-stage semantic parsing. *arXiv preprint arXiv:2005.13485*.
- Dheeru Dua, Yizhong Wang, Pradeep Dasigi, Gabriel Stanovsky, Sameer Singh, and Matt Gardner. 2019. Drop: A reading comprehension benchmark requiring discrete reasoning over paragraphs. *arXiv preprint arXiv:1903.00161*.
- Charles T Hemphill, John J Godfrey, and George R Doddington. 1990. The atis spoken language systems pilot corpus. In *Speech and Natural Language: Proceedings of a Workshop Held at Hidden Valley, Pennsylvania, June 24-27, 1990*.
- Matthew Honnibal, Ines Montani, Sofie Van Lan-deghem, and Adriane Boyd. 2020. *spaCy: Industrial-strength Natural Language Processing in Python*.
- Aishwarya Kamath and Rajarshi Das. 2018. A survey on semantic parsing. *arXiv preprint arXiv:1812.00978*.
- Anjuli Kannan, Karol Kurach, Sujith Ravi, Tobias Kaufmann, Andrew Tomkins, Balint Miklos, Greg Corrado, Laszlo Lukacs, Marina Ganea, Peter Young, et al. 2016. Smart reply: Automated response suggestion for email. In *Proceedings of the 22nd ACM SIGKDD International Conference on Knowledge Discovery and Data Mining*, pages 955–964.
- Younes Karimi, Anna Squicciarini, and Shomir Wilson. 2022. Automated detection of doxing on twitter. *arXiv preprint arXiv:2202.00879*.
- Sungdong Kim, Minsuk Chang, and Sang-Woo Lee. 2021. Neuralwoz: Learning to collect task-oriented dialogue via model-based simulation. *arXiv preprint arXiv:2105.14454*.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Ves Stoyanov, and Luke Zettlemoyer. 2019. Bart: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension. *arXiv preprint arXiv:1910.13461*.
- Ildikó Pilán, Pierre Lison, Lilja Øvrelid, Anthi Papadopoulou, David Sánchez, and Montserrat Batet. 2022. The text anonymization benchmark (tab): A dedicated corpus and evaluation framework for text anonymization. *arXiv preprint arXiv:2202.00443*.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. 2019. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9.
- Timo Schick and Hinrich Schütze. 2021. Generating datasets with pretrained language models. *arXiv preprint arXiv:2104.07540*.
- Nathan Schucher, Siva Reddy, and Harm de Vries. 2021. The power of prompt tuning for low-resource semantic parsing. *arXiv preprint arXiv:2110.08525*.
- Richard Shin, Christopher H Lin, Sam Thomson, Charles Chen, Subhro Roy, Emmanouil Antonios Platanios, Adam Pauls, Dan Klein, Jason Eisner, and Benjamin Van Durme. 2021. Constrained language models yield few-shot semantic parsers. *arXiv preprint arXiv:2104.08768*.
- Alane Suhr, Stephanie Zhou, Ally Zhang, Iris Zhang, Huajun Bai, and Yoav Artzi. 2018. A corpus for reasoning about natural language grounded in photographs. *arXiv preprint arXiv:1811.00491*.

Amazon Transcribe. [Amazon Transcribe now supports redaction of personal identifiable information \(PII\) for streaming transcriptions](#). Accessed: 2022-03-08.

Bo-Hsiang Tseng, Yinpei Dai, Florian Kreyszig, and Bill Byrne. 2021. Transferable dialogue systems and user simulators. *arXiv preprint arXiv:2107.11904*.

Yushi Wang, Jonathan Berant, and Percy Liang. 2015. Building a semantic parser overnight. In *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)*, pages 1332–1342.

Tomer Wolfson, Mor Geva, Ankit Gupta, Matt Gardner, Yoav Goldberg, Daniel Deutch, and Jonathan Berant. 2020. Break it down: A question understanding benchmark. *Transactions of the Association for Computational Linguistics*, 8:183–198.

Shan Wu, Bo Chen, Chunlei Xin, Xianpei Han, Le Sun, Weipeng Zhang, Jiansong Chen, Fan Yang, and Xunliang Cai. 2021. From paraphrasing to semantic parsing: Unsupervised semantic parsing via synchronous semantic decoding. *arXiv preprint arXiv:2106.06228*.

Chunyang Xiao, Marc Dymetman, and Claire Gardent. 2016. Sequence-based structured prediction for semantic parsing. In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1341–1350.

Pengcheng Yin, John Wieting, Avirup Sil, and Graham Neubig. 2021. On the ingredients of an effective zero-shot semantic parser. *arXiv preprint arXiv:2110.08381*.

Zijian Zhao, Su Zhu, and Kai Yu. 2019. Data augmentation with atomic templates for spoken language understanding. *arXiv preprint arXiv:1908.10770*.

Victor Zhong, Mike Lewis, Sida I Wang, and Luke Zettlemoyer. 2020. Grounded adaptation for zero-shot executable semantic parsing. *arXiv preprint arXiv:2009.07396*.

A Masking and Replacing Personal Identifiable Information

A.1 ATIS

The ATIS grammar is somewhat loosely defined and does not clearly indicate the instances of PII. This would be problematic in a real production setting due to making it difficult to guarantee masking out all PII. However, for our experiments we simply truecase the data and apply named entity recognition using spaCy (Honnibal et al., 2020), which we find is highly successful from a qualitative inspection. We treat detected named entities as PII.

To remove PII, we devise two methods: 1. masking and 2. generating entirely new plans. In 1., we apply the above method to detect PII, mask it with the entity type, and ask GPT3 to infill (Figure 2). In 2., corresponding to our GPT-RANK method, we feed GPT3 example plans from the seed data and ask for an entirely new plan that does not contain PII (Figure 3).

Truecase	NER recognition + mask	GPT3 fill based on entity type
<ul style="list-style-type: none">• "Show me all the flights arriving in Charlotte around next Monday please "	<ul style="list-style-type: none">• "Show me all the flights arriving in [GPE] around [TIME] please "	<ul style="list-style-type: none">• "Show me all the flights arriving in Austin around 5 pm please "

Figure 2: An example of infill masking PII.

A.2 SMCaFlow

Since the SMCaFlow grammar (described in detail in Appendix A.2 of Shin et al. (2021)) is type-annotated, we define three categories of PII: names, event titles, and locations. Each category is easily identifiable from the logical form, so it suffices to sample a new value from the same category in the logical form to guarantee that PII is replaced.

We sample names from a distribution balanced for ethnicity and gender. For event titles and locations, we sample them from GPT3 by prompting with seed data canonical forms containing event titles and/or locations, and then prefixing the generation with `find event called something like "` (event titles) or `a mix of weather at "` and `find event at "` (locations). We cut off the generation once the next `"` appears.

INPUT: return ground transportation ;return #1 in boston ;return #2 between the airport ;return #3 and downtown

return flights ;return #1 that are nonstop ;return #2 from san diego ;return #3 to new york

return flights ;return #1 from denver ;return #2 to pittsburgh ;return #3 leaving after 6pm ;return #3 leaving before 7pm ;return #1 of both #4 and #5

GPT3: return flights ;return #1 that are nonstop ;return #2 between denver ;return #3 and oakland ;return #4 leaving after noon ;return #5 arriving after 5pm

Figure 3: INPUT is an example prompt, consisting of several plans from the seed data, given to GPT3 to generate new plans to begin the data augmentation procedure in GPT-RANK.

B Example Parses

We include some additional example parses for USER-RANK-3X compared to BASE on SMCaFlow, as well as GPT-RANK-3X compared to BASE on ATIS, when both start with the same seed data (Tables 7, 8, 9).

We additionally provide some qualitative analysis on the SMCaFlow examples. Most examples that both methods get correct are relatively simple intents (Example 1). However, there are also many examples where one or both methods output a parse which a human might judge semantically equivalent to the true parse (Examples 2 and 3), suggesting that our exact match metric might underrepresent the true strength of the models.

The next few examples illustrate cases in which USER-RANK-3X improves over BASE, for example by being better at segmenting names (Examples 4). BASE is also sometimes just wildly incorrect on examples where USER-RANK-3X matches the gold exactly, or at least closely in meaning (Examples 5 and 6). Of course, given the tiny amount of initial training data, USER-RANK-3X is not immune to these types of errors either (Example 7).

Example 1	
Previous Agent	<i>I've deleted your event matching "doctor's appointment".</i>
Natural	<i>thank you</i>
BASE	Thanks for your help!
USER-RANK-3X	Thanks for your help!
Gold Canonical	Thanks for your help!
Example 2	
Previous Agent	<i>Ok, I won't do anything for now. Let me know what else I can help you with.</i>
Natural	<i>Make an appointment for tomorrow from 1pm to 6pm</i>
BASE	create event starting tomorrow 1 PM to 6 PM
USER-RANK-3X	create event starting tomorrow 1 PM to 6 PM
Gold Canonical	create event starting tomorrow 1 PM ending 6 PM after that datetime
Example 3	
Previous Agent	<i>Next is your tour potential development sites on Thursday the 28th from 3:00 to 5:00 PM.</i>
Natural	<i>I want to delete that one.</i>
BASE	delete the event
USER-RANK-3X	delete find event
Gold Canonical	delete the event
Example 4	
Previous Agent	<i>Does one of these work?</i>
Natural	<i>When is Easter next year?</i>
BASE	start time of find event called something like "Easter next year"
USER-RANK-3X	start time of find event called something like "Easter" starting next year
Gold Canonical	Easter next year

Table 7: Example parses by the baseline BASE and our best method USER-RANK-3X on SMCaFlow.

Example 5	
Previous Agent	
Natural	<i>list to me my calendar please</i>
BASE	create event on today afternoon
USER-RANK-3X	find event
Gold Canonical	find event
Example 6	
Previous Agent	<i>The "library" is on Monday the 30th from 10:00 to 10:30 AM.</i>
Natural	<i>Ok! now tell me when does my Coffee Date start?</i>
BASE	ERROR: can't answer trivia
USER-RANK-3X	start time of find event called something like "Coffee Date"
Gold Canonical	start time of find event called something like "Coffee Date"
Example 7	
Previous Agent	
Natural	<i>Hey, I was wondering who the organizer is for the museum event next week.</i>
BASE	ERROR: can't answer trivia
USER-RANK-3X	ERROR: can't answer trivia
Gold Canonical	organizer of find event called something like "museum" during next week

Table 8: Additional example parses by the baseline BASE and our best method USER-RANK-3X on SMCaFlow.

Example 8	
Previous Agent	
Natural	<i>I want a flight from houston to memphis on tuesday morning</i>
BASE	return flights ;return #1 from houston ;return #2 to memphis ;return #3 on tuesday morning
GPT-RANK-3X	return flights ;return #1 from houston ;return #2 to memphis ;return #3 on tuesday morning
Gold Canonical	return flights ;return #1 from houston ;return #2 to memphis ;return #3 on tuesday ;return #4 in the morning
Example 9	
Previous Agent	
Natural	<i>What ground transportation is available from the pittsburgh airport to downtown and how much does it cost</i>
BASE	return transportation ;return #1 that is ground ;return #2 to downtown pittsburgh; return cost of #4
GPT-RANK-3X	return transportation ;return #1 that is ground ;return #2 from the pittsburgh ;return #3 to downtown pittsburgh; return cost #4
Gold Canonical	return ground transportation ;return #1 which is available ;return #2 from the pittsburgh airport ;return #3 to downtown ;return the cost of #4

Table 9: Additional example parses by the baseline BASE and our best method GPT-RANK-3X on ATIS.

C Preliminary Experiments on Other Break Subsets

We additionally ran preliminary experiments on the DROP (Dua et al., 2019) (reading comprehension) and NLVR2 (Suhr et al., 2018) (language-vision reasoning) subsets of Break (Wolfson et al., 2020). We used a similar setup to our ATIS and SMCaFlow experiments, with 30 initial seed data \mathcal{D} and 300 unlabeled user utterances \mathcal{U} .

However, across multiple trials of multiple methods (BASE, USER-RANK, GPT-RANK, USER-CYC), we never observed performance above 2% on either domain. This may be partially due to the diversity of the data; for example, DROP is an amalgamation of data from several sources. However, we hypothesize that this across-the-board poor performance is primarily the result of an SCFG for canonical utterances which results in somewhat unnatural language (Table 10), and that performance could be greatly improved with a better SCFG. Given the current form of our canonical utterances in DROP and NLVR2, it is challenging to learn the task given just 30 seed examples. In comparison, the SMCaFlow canonical utterances (Table 1 in the main text) are much more natural.

DROP Natural	<i>Which player had the shortest touchdown reception of the game?</i>
DROP Canonical	return touchdown receptions ;return shortest of #1 ;return player of #2
NLVR2 Natural	<i>If there are two carts, but only one of them has a canopy.</i>
NLVR2 Canonical	return carts ;return number of #1 ;return if #2 is equal to two ;return canopy ;return #1 that has #4 ;return number of #5 ;return if #6 is equal to one ;return if both #3 and #7 are true

Table 10: Examples of natural utterances with corresponding canonical utterances for DROP and NLVR2 domains. The language of the canonical utterances is relatively unnatural.

We additionally inspect some inaccurate example predictions by BASE on DROP and NLVR2, which are often wildly incorrect (Table 11). We also show some example (c, n) pairs generated by our data augmentation procedure, demonstrating the failure to propose good natural language n given the limited data and unnatural canonical c (Table 12).

DROP Natural	<i>Which player threw more yards in the game, Young or Manning?</i>
DROP Top-1 Parse	return that was the highest ;return that was more of #1 ;return number of #2 for each #1 ;return #1 where #3 is lower than one ;return number of #4
NLVR2 Natural	<i>If there are bananas with stickers on them</i>
NLVR2 Top-1 Parse	return, ;return number of #1 ;return if #2 is equal to one

Table 11: Predictions by BASE on DROP and NLVR2 which are wildly incorrect. Our data augmentation methods fare no better.

DROP Canonical	return the five
DROP Simulated Natural	Fact-checkers failed to catch five factual errors.
NLVR2 Canonical	return left image ;return #1 that are dirty ;return if #2 is in one of the images
NLVR2 Simulated Natural	If any of the trucks are dirty.

Table 12: Example simulated natural utterances generated by prompting GPT3 on DROP and NLVR2, after reranking and selecting the best of 20 generations. The correspondence between canonical and simulated natural utterances remains imperfect.