

---

# Grammar Prompting for Domain-Specific Language Generation with Large Language Models

---

Bailin Wang<sup>◇</sup> Zi Wang<sup>†</sup> Xuezhi Wang<sup>†</sup> Yuan Cao<sup>‡</sup> Rif A. Saurous<sup>†</sup> Yoon Kim<sup>◇</sup>  
<sup>◇</sup>Massachusetts Institute of Technology <sup>†</sup>Google DeepMind <sup>‡</sup>Google Research  
{bailinw, yoonkim}@mit.edu, {wangzi, xuezhiw, yuanc, rif}@google.com

## Abstract

Large language models (LLMs) can learn to perform a wide range of natural language tasks from just a handful of in-context examples. However, for generating strings from highly structured languages (e.g., semantic parsing to complex domain-specific languages), it is challenging for the LLM to generalize from just a few exemplars. We propose *grammar prompting*, a simple approach to enable LLMs to use external knowledge and domain-specific constraints, expressed through a grammar in Backus–Naur Form (BNF), during in-context learning. Grammar prompting augments each demonstration example with a specialized grammar that is minimally sufficient for generating the particular output example, where the specialized grammar is a subset of the full DSL grammar. For inference, the LLM first predicts a BNF grammar given a test input, and then generates the output according to the rules of the grammar. Experiments demonstrate that grammar prompting can enable LLMs to perform competitively on a diverse set of DSL generation tasks, including semantic parsing (SMCalFlow, Overnight, GeoQuery), PDDL planning, and SMILES-based molecule generation.

## 1 Introduction

Prompting large language models (LLMs) with demonstrations optionally combined with natural language instructions has been shown to be an effective approach for surfacing their myriad capabilities acquired through pretraining [10]. This approach is however inadequate for applications where the task specifications cannot be fully delineated through just a handful of exemplars, for example in semantic parsing where an LLM must translate a natural language utterance to an executable program in a domain-specific language (DSL). DSLs often incorporate domain-specific abstractions and semantics that are difficult to characterize via just a few demonstrations. And unlike general-purpose programming languages, DSLs are by definition specialized and thus unlikely to have been encountered often enough (or at all) during pretraining for the LLM to acquire its full syntax.

How can we draw on the few-shot learning capabilities of LLMs to generate structured strings that are substantially different from those seen during pretraining? This work explores *grammar prompting* as a simple approach for data-efficient generation of structured languages where an output string in the language can be derived through a series of symbolic manipulations. We exploit the fact that constraints over a structured output space can often be succinctly described by a context-free grammar in Backus–Naur Form (BNF), which is commonly used to define the syntax of a language. Grammar prompting augments each in-context example  $(x, y)$  with a *specialized* BNF grammar  $G[y]$  that is minimally sufficient for generating  $y$ . Given a new input, the LLM first predicts the specialized BNF grammar and then generates the answer conditioned on the grammar.

Grammar prompting follows the recent line of work which enhances the few-shot reasoning capabilities of LLMs by interleaving intermediate “reasoning” steps between each in-context input and

Code and data available at: <https://github.com/berlino/grammar-prompting>.

output [51, 24, 86, 80, 73]. The key difference in our approach is that the intermediate variable is in the form of a formal grammar rather than in natural language, which focuses on eliciting the symbolic manipulation capabilities of LLMs. The use of a formal grammar moreover makes it possible to impose constraints during incremental decoding such that syntactic validity is guaranteed. Finally, unlike chain-of-thought-style prompts [86] which typically require manual verbalization of the intermediate reasoning steps, in our approach the specialized grammar  $G[\mathbf{y}]$  can be derived automatically by parsing the output  $\mathbf{y}$  with the full (unspecialized) DSL grammar.

To summarize,

- We propose grammar prompting as a simple approach for enabling LLMs to generate highly structured languages from just a few exemplars.
- We design a constrained LLM decoding algorithm tailored to grammar prompting, which guarantees syntactic validity while minimizing the number of LLM API calls.
- We apply grammar prompting to various domain specific languages for semantic parsing (SMCalFlow, Overnight, GeoQuery), AI planning (PDDL), and molecule generation (SMILES), and find that it can meaningfully improve upon standard prompting baselines in the few-shot setting.

## 2 Background

In this section, we define our problem and review the few-shot learning method that we build on.

### 2.1 Problem Formulation: Domain-Specific Language Generation

Let  $\Sigma^*$  be the set of all finite strings over an alphabet  $\Sigma$ , and further let  $D \subseteq \Sigma^*$  be a domain-specific language (DSL) for an application of interest. Given an input  $\mathbf{x}$  (e.g., a natural language command) we are interested in generating  $\mathbf{y} \in D$  (e.g., a program in a DSL fulfilling the command), as shown by the following calendar assistant example from SMCaFlow [6]:

$\mathbf{x}$  : Add meeting with Jean’s manager on Wednesday at 3PM.  
 $\mathbf{y}$  : CreateEvent(&(start\_? WednesdayNumberPM(3))(attendee\_? FindManager(Jean)))

DSLs are crafted by experts who use their domain-specific knowledge to incorporate higher-level abstractions than are typically found in general-purpose programming languages. We assume access to an expert-defined grammar  $G$  that fully specifies the DSL’s syntax. As is the case with many DSLs, we further assume that  $G$  is a context-free grammar in Backus–Naur Form (BNF). See Figure 1 for a simple example adapted from SMCaFlow [6]. Letting  $L(G)$  be the language generated by  $G$ , we have  $D \subseteq L(G) \subseteq \Sigma^*$  (not all syntactically valid programs are semantically valid).

event	::=	"CreateEvent(" constraint ")"
		"QueryEvent(" constraint ")"
constraint	::=	"(& constraint constraint)"
		"(start_? date time?)"
		"(attendee_? attendee*)"
date	::=	"Wednesday"   "Monday"
number	::=	("0".."9")+
time	::=	"NumberAM(" number ")"
		"NumberPM(" number ")"
attendee	::=	"Bob"   "Carol"   "Jean"
		"FindManager"

Figure 1: A simple BNF grammar for a calendar DSL.

### 2.2 Few-shot Learning with Large Language Models

In-context learning with large language models (LLMs) has been shown to be an effective approach for few-shot learning [10]. Under this approach, a pretrained LLM is conditioned on  $N$  demonstration examples  $(\mathbf{x}^{(i)}, \mathbf{y}^{(i)})_{i=1}^N$  followed by a test example  $\mathbf{x}$ , and the output is given by decoding from the prompted LLM, i.e.,  $P_{\text{LLM}}(\mathbf{y} | \mathbf{x}, (\mathbf{x}^{(i)}, \mathbf{y}^{(i)})_{i=1}^N)$ . The demonstration examples can be optionally preceded by natural language instructions to further improve performance or even enable zero-shot learning [85, 62]. Recent work has additionally shown that interleaving natural language verbalizations of intermediate reasoning steps between each  $\mathbf{x}^{(i)}$  and  $\mathbf{y}^{(i)}$  can greatly improve few-shot performance on complex reasoning tasks [51, 86, 80, 73, 16].

The effectiveness of few-shot in-context learning depends both on how useful the implicit knowledge acquired through pretraining is for the task, and on how effectively the task specifications can be conveyed through the demonstrations. For DSL, the structured nature of combinatorial output space (i.e., the DSL grammar  $G$ ) cannot be adequately captured through just a handful of demonstrations. Thus, few-shot generation of strings of a DSL remains challenging for LLMs.

LLM Prompt	
You are an expert programmer, and you need to write a program for the given natural language query. First, you should write a grammar that contains all the necessary BNF rules. Then, you should write programs that conform to your predicted rules.	
(optional) $G$ :	[BEGIN RULES]                    ...                    [END RULES]
$x^{(1)}$ :	find the meeting on Wednesday with Bob and Carol
$G[y^{(1)}]$ :	<pre> event      ::= "QueryEvent(" constraint ")" constraint ::= "&amp;" constraint constraint ")"              "(start_?" date ")"              "(attendee_?" attendee attendee ")" date       ::= "Wednesday" attendee   ::= "Bob"   "Carol" </pre>
$y^{(1)}$ :	QueryEvent(& (start_? Wednesday)(attendee_? Bob Carol))
$x$ :	Add meeting with Jean's manager on Wednesday at 3PM
LLM Output	
$\hat{G}$ :	<pre> event      ::= "CreateEvent(" constraint ")" constraint ::= "&amp;" constraint constraint ")"              "(start_?" date time ")"              "(attendee_?" attendee ")" date       ::= "Wednesday" time       ::= "NumberPM(3)" attendee   ::= "FindManager(" attendee ")"   "Jean" </pre>
$\hat{y}$ :	CreateEvent(& (start_? Wednesday NumberPM(3))(attendee_? FindManager(Jean)))

**Figure 2:** Example of grammar prompting for a calendar DSL. We interleave the minimal specialized grammar  $G[y^{(i)}]$  between the demonstrations  $x^{(i)}$  and  $y^{(i)}$ . During decoding, the LLM first predicts the specialized grammar  $\hat{G}$ , and then predicts the program  $\hat{y}$  conditioned on  $\hat{G}$ . The blue portion is not part of the actual prompt and only shown for illustrative purposes.

### 3 Grammar Prompting

Grammar prompting exploits the fact that while the actual strings of a DSL may not have been encountered frequently enough (or at all) during pretraining for the LLM to implicitly acquire its syntax, the LLM will likely have encountered many instances of *metalanguages* (languages used to describe other languages). BNF grammars are a standard metalanguage for specifying a language’s syntax, and are expected to occur in the LLM training corpus with some frequency (e.g., in computer science textbooks). We thus focus on using BNF grammars for few-shot DSL generation.

Let  $G = \bigcup_{j=1}^M \{r_j\}$  be an extended BNF grammar where each rule  $r_j$  is of the form

$$\langle \text{symbol} \rangle ::= \langle \text{expr}_1 \rangle \mid \langle \text{expr}_2 \rangle \mid \dots$$

Here  $\langle \text{symbol} \rangle$  is a nonterminal symbol and each  $\langle \text{expr}_1 \rangle$  is a sequence of nonterminal and terminal symbols.<sup>1</sup> A straightforward approach for incorporating a BNF grammar during in-context learning is to simply prepend the string representation of the full grammar  $G$  to the demonstration examples, along with an instruction to use the grammar. However in preliminary experiments, we found that this did not yield any improvements.<sup>2</sup>

#### 3.1 Specialized Grammars

We propose to use *specialized grammars* to enable better use of domain-specific knowledge and constraints. A specialized grammar  $G' \subseteq G$  is a grammar obtained from taking a subset of the rules of the full grammar  $G$ . We further define  $G[y]$ , a *minimal specialized grammar* of  $y$ , to be a BNF

<sup>1</sup>For brevity we forgo the formal tuple-based definition of  $G$  and instead define  $G$  to be equivalent to its context-free rules. We also freely go back and forth between this set definition of  $G$  and its string representation.

<sup>2</sup>However, when combined with specialized grammars we did observe small improvements by appending the full DSL grammar to the instructions. Hence, for all experiments where  $G$  is small enough (GeoQuery, Overnight-B, SMILES), we include  $G$  as part of the instruction. See Figure 2.

grammar with the following properties: (1)  $\mathbf{y} \in L(G[\mathbf{y}])$ , and (2)  $\forall r \in G[\mathbf{y}], \mathbf{y} \notin L(G[\mathbf{y}] \setminus \{r\})$ .<sup>3</sup> We can readily obtain a minimal specialized grammar by using  $G$  to parse  $\mathbf{y}$  and then taking the union of rules that were used in the derivation of  $\mathbf{y}$ .

Grammar prompting feeds a sequence of  $(\mathbf{x}^{(i)}, G[\mathbf{y}^{(i)}], \mathbf{y}^{(i)})_{i=1}^N$  along with  $\mathbf{x}$  as a prompt to an LLM. For inference we first obtain the specialized grammar with an (approximate) arg max decoding

$$\hat{G} = \arg \max_{G' \subseteq G} P_{\text{LLM}}(G' | \mathbf{x}, (\mathbf{x}^{(i)}, G[\mathbf{y}^{(i)}], \mathbf{y}^{(i)})_{i=1}^N).$$

We then obtain the program conditioned on  $\hat{G}$ ,

$$\hat{\mathbf{y}} = \arg \max_{\mathbf{y} \in L(\hat{G})} P_{\text{LLM}}(\mathbf{y} | \hat{G}, \mathbf{x}, (\mathbf{x}^{(i)}, G[\mathbf{y}^{(i)}], \mathbf{y}^{(i)})_{i=1}^N).$$

We discuss how to perform constrained decoding with  $\hat{G} \subseteq G$  and  $\hat{\mathbf{y}} \in L(\hat{G})$  in the next section. Grammar prompting views DSL program generation as a *grammar specialization* process where given a natural language specification  $\mathbf{x}$ , a set of production rules,  $\hat{G}$ , is selected from  $G$ , and then a program  $\hat{\mathbf{y}}$  is deduced according to the selected rules. Grammar prompting can also be viewed as an instance of chain-of-thought prompting [51, 86] where the intermediate thought is in the form of a formal grammar. However, unlike typical chain-of-thought prompting where the answer is (usually) deterministic given the intermediate reasoning steps, in our case there is still some uncertainty with respect to  $\hat{\mathbf{y}}$  given  $\hat{G}$  (e.g.,  $L(\hat{G})$  could still be infinite).

### 3.2 Constrained Decoding

The use of a formal grammar as an intermediate variable makes it possible to enforce grammatical constraints during autoregressive LLM decoding. We first discuss how we enforce the constraint  $\mathbf{y} \in L(\hat{G})$ . One approach to constrained decoding is to use  $\hat{G}$  to obtain a left-to-right Earley parser [18] and only decode from valid continuations at each decoding step. However this simple strategy poses several practical challenges when working with API-only LLMs. For one, a valid terminal continuation in  $\hat{G}$  may consist of multiple BPE tokens. Moreover, while we can sample a valid continuation at each time step by disallowing invalid tokens,<sup>4</sup> since the set of valid continuations changes at each time step, this strategy would require calling the LLM API at each time step with the full prompt and prefix along with the disallowed continuations, which is prohibitively expensive.<sup>5</sup>

While there are many methods for grammar-constrained LM decoding [68, 64, 26], we present a simple strategy which speculatively decodes from the LLM to look ahead for multiple tokens. The pseudocode is shown in Algorithm 1. At each prediction step, we ask the LLM to speculatively decode the full program conditioned on the current prefix (lines 4-5). If the resulting continuation leads to a valid program, we return it (lines 6-7). Otherwise, we consult an Earley parser to extract the longest valid prefix from the current prediction ( $\mathbf{y}_{\text{prefix}}$ ), along with a set of valid terminals that can

---

#### Algorithm 1 Earley-based Constrained Generation

---

**Input:** Test input  $\mathbf{x}$ , predicted grammar  $\hat{G}$

**Output:** Program  $\hat{\mathbf{y}} \in L(\hat{G})$

---

```

1:  $\hat{\mathbf{y}} \leftarrow \epsilon$  ▷ initialize to empty string
2: while True do
3:    $\bar{\mathbf{y}} \leftarrow \text{decode} \left( P_{\text{LLM}}(\cdot | \mathbf{x}, \hat{G}, \hat{\mathbf{y}}, \dots) \right)$ 
4:    $\hat{\mathbf{y}} \leftarrow \hat{\mathbf{y}} \cdot \bar{\mathbf{y}}$  ▷ concatenation
5:   if  $\hat{\mathbf{y}} \in L(\hat{G})$  then ▷ try parsing with  $\hat{G}$ 
6:     return  $\hat{\mathbf{y}}$  ▷ return if successful
7:   else ▷ if parsing fails, need to correct
8:      $\mathbf{y}_{\text{prefix}}, \Sigma[\mathbf{y}_{\text{prefix}}] \leftarrow \text{EarleyParse}(\hat{\mathbf{y}}, \hat{G})$ 
9:      $\mathbf{w}^* \leftarrow \arg \max_{\mathbf{w} \in \Sigma[\mathbf{y}_{\text{prefix}}]} P_{\text{LLM}}(\mathbf{w} | \mathbf{y}_{\text{prefix}}, \dots)$ 
10:     $\hat{\mathbf{y}} \leftarrow \mathbf{y}_{\text{prefix}} \cdot \mathbf{w}^*$ 
11:   end if
12: end while
13: return  $\hat{\mathbf{y}}$ 

```

---

<sup>3</sup>Note that  $\mathbf{y}$  may have more than one minimal specialized grammar due to the potential instantiation of extended BNF rules. For instance, the expression "(attendee\_?" attendee+ ")" depicted in Figure 1 implicitly defines all occurrences of attendee greater than 1. If this expression is incorporated into a program, either the concrete rule "(attendee\_?" attendee ")" or the original rule could be included in the minimal specialized grammar. In most applications we consider the rules of the minimal specialized grammar will be concrete, and thus there will only be one parse associated with  $\mathbf{y}$ . See appendix A.1 for further details.

<sup>4</sup>For example by using the `logit_bias` argument from OpenAI's LLM API.

<sup>5</sup>These costs might be mitigated in the future if LLM APIs allow for cheaper use of cached prompts.

Earley-based Constrained Decoding
$\hat{y}^{(t)} : \text{CreateEvent}(\& (\text{start\_? Wednesday NumberPM}(3))(\text{attendee\_? Jean's Manager}))$ <hr/> <i>Extract the longest valid prefix, and possible fixes (i.e., next terminals) based on Earley parsing:</i> $y_{\text{prefix}} : \text{QueryEvent}(\& (\text{start\_? Wednesday})(\text{attendee\_?}$ $\Sigma[y_{\text{prefix}}]: \{ \text{Jean}, \text{FindManager} \}$ <hr/> <i>Find the best candidate and concatenate it with the prefix:</i> $\hat{y}_{(t)} \leftarrow \text{QueryEvent}(\& (\text{start\_? Wednesday})(\text{attendee\_? FindManager}(\$

**Figure 3:** Illustration of how an predicted program is corrected in our proposed Earley-based constrained decoding. The final partial program will be subsequently fed into the LLM for continuation.

follow the prefix ( $\Sigma[y_{\text{prefix}}]$ ). Finally, we rely on the LLM’s probabilities to decide which terminal to use, with which a new partial program can be constructed (lines 10-11).<sup>6</sup> Figure 3 illustrates one prediction step where the predicted program is corrected into a new valid partial program. Note that  $w$  can consist of multiple BPE tokens, e.g., "FindManager(" in Figure 3. By scoring over multi-token terminals, the search procedure is implicitly augmented by looking ahead for a few tokens.

We use a similar procedure to operationalize the constraint  $G' \subseteq G$ , except that EarleyParse (used at Algorithm 1, line 9) is constructed with a *metagrammar* (i.e., the grammar of  $G$ ) for grammar prediction. See appendix A.1 for more details. In our ablation study we find that while these constraints are helpful insofar as they guarantee syntactic validity, grammar prompting still meaningfully improves upon standard prompting with even with simple unconstrained decoding.

## 4 Experiments

We apply grammar prompting to diverse domains: DSLs for semantic parsing (SMCalFlow, Overnight, GeoQuery), an action DSL (PDDL planning), and a molecule generation DSL (SMILES). These experiments are not necessarily intended to improve upon the state-of-the-art on these benchmarks but rather intended to assess whether LLMs can improve upon standard prompting for few-shot DSL generation by learning to predict and use grammars during in-context learning.

### 4.1 Semantic Parsing for Tool Usage

Software tools are typically accompanied by a collection of human-interpretable APIs which provide a platform for developers to interact programmatically with the tools. These APIs constitute a DSL, where each production rule of the grammar specifies the input and output types for a specific API call (see Figure 1 for an example). These tools demonstrate a broad spectrum in terms of DSL complexity, ranging from single-function tools such as Google(user\_query), Translate(sentence, language) to more complex tools such as the entirety of Wolfram language.<sup>7</sup> Enabling LLMs to use external tools via APIs is an important step towards enhancing their capabilities [63, 56, 53, 72, 47].

We test our approach on standard semantic parsing benchmarks involving complex DSLs: SM-CalFlow [6], which features human-generated utterances about calendar management (see Figure 2); GeoQuery [99] which features queries against a US Geography database; and Overnight-Blocks [81], which features queries about blocks in a synthetic block world. See appendix B for examples of input-output pairs along with the specialized grammars. The original benchmarks target the training of conventional semantic parsers and thus contain hundreds/thousands of training examples. Even prompting-based approaches on these benchmark rely on retrieval-based in-context learning which first retrieves  $m$  exemplars from a large training set of  $n$  examples ( $n \gg m$ ) based on some similarity measure (e.g., BM-25), and then performs in-context learning with the retrieved exemplars [57, 95, 68, 46]. In contrast, we target the true few-shot setting where we only assume access to 16–32 demonstration examples.

Our baselines here include: (1) standard prompting, (2) standard prompting with constrained decoding based on the full DSL grammar  $G$  [68, 64], and (3) a derivation tree-based prompting baseline which imbues more structural information to the exemplars by feeding the linearized derivation

<sup>6</sup>In rare cases the set  $\Sigma[y_{\text{prefix}}]$  was too large to feed to LLM APIs. In these cases we used Sentence-BERT [59] to compute the similarity between the encoding of  $y_{\text{prefix}} \cdot w$  and  $\hat{y}^{(t)}$  and took the top 16 candidates as  $\Sigma[y_{\text{prefix}}]$ .

<sup>7</sup><https://www.wolfram.com/language/>

Approach	GeoQuery		SMCalFlow	Overnight-Blk	
	Prog.	Exec.	Prog.	Prog.	Exec.
Standard Prompting (unconstrained decoding)	60.7	81.5	46.4	29.3	54.7
w. <i>constrained decoding</i> ( $\hat{\mathbf{y}} \in L(G)$ )	61.1	81.8	49.2	29.3	54.7
Linearized Derivation Tree Prompting	58.6	77.5	50.0	27.3	56.4
Grammar Prompting (unconstrained decoding)	67.1	87.5	50.8	34.8	57.4
w. <i>grammar constraint</i> ( $\hat{G} \subseteq G$ )	67.9	88.6	51.3	37.1	60.4
w. <i>grammar and program constraint</i> ( $\hat{\mathbf{y}} \in L(\hat{G})$ )	69.6	88.9	52.4	37.6	60.9
w. <i>oracle grammar</i> ( $\hat{G} = G[\mathbf{y}]$ )	95.7	96.1	80.0	73.9	94.2
w. <i>oracle grammar + program constraint</i>	95.7	96.8	83.6	74.4	96.5

**Table 1:** Results on few-shot semantic parsing with Codex with various decoding strategies. GeoQuery and Overnight-Blk use 32 in-context examples, and SMCalFlow uses 16 examples. We show both program (Prog.) and execution (Exec.) accuracy when possible.

Model	Retrieval-based ICL			GeoQuery Out-of-Distribution			
	GeoQuery	SMCalFlow	Overnight-Blk	Template	TMCD	Length	NewFunc
(# ICL examples / # retrieval set)	(32/560)	(16/128)	(32/1,436)	(32/441)	(32/440)	(32/440)	(32/453)
Previous Work	86.1 <sup>♣</sup>	60.7 <sup>♣</sup>	65.2 <sup>◇</sup>	–	–	–	–
Standard Prompting	96.8	60.0	69.4	93.2	77.1	86.4	63.3
Grammar Prompting	97.9	62.8	70.2	95.7	86.6	88.6	90.8
w. <i>oracle grammar</i>	98.6	88.9	97.2	97.9	95.0	95.7	96.2

**Table 2:** Results on retrieval-based in-context learning (left) and compositional generalization (right) with Codex. GeoQuery and Overnight-Blk show execution accuracy while SMCalFlow shows program accuracy. The numbers with <sup>♣</sup>, <sup>♠</sup> and <sup>◇</sup> are taken from Herzig and Berant [31], Ye et al. [95] and Cao et al. [11], respectively.

tree instead of the surface form program.<sup>8</sup> We use Codex-davinci-002 [13] as the base LLM for these main experiments. Language models trained on code (such as Codex) have shown to be particularly effective on semantic parsing benchmarks [67]. We evaluate according to program accuracy (matching the predicted and reference programs) as well as execution accuracy (same execution in both programs) if possible.

**Few-shot results.** The main results are shown in Table 1. We find that grammar prompting can meaningfully improve upon the standard prompting baseline even without constrained decoding. Interestingly, grammar prompting outperforms derivation tree prompting which actually provides *more* information than the minimal specialized grammar  $G[\mathbf{y}]$  (since the derivation tree explicitly shows how the rules are actually applied to obtain the program). This potentially indicates that having the LLM “plan out” the program by forcing it to predict the specialized grammar  $\hat{G}$  first is an effective strategy. We also analyze the effect of constrained decoding on the number of LLM API calls in Table 7 of appendix A.1, where we observe that constrained decoding requires roughly three times more API calls than unconstrained decoding. However, despite the promising performance of grammar prompting, there is a large gap between using the predicted grammar versus using the oracle grammar (i.e., setting  $\hat{G} = G[\mathbf{y}]$ ), indicating opportunities for further work in this area.

**Retrieval-based in-context learning.** While our core target application is few-shot semantic parsing, we also apply grammar prompting for retrieval-based in-context learning to test whether it can still improve performance in the data-rich regime and also to compare against prior work on these benchmarks. Results in Table 2 (left) show that grammar prompting can improve results even in this setting, although the improvements are less pronounced than in the few-shot setting.

**Out-of-distribution generalization.** We experiment to see whether grammar prompting can improve compositional generalization on GeoQuery. Specifically, we test grammar prompting on the compositional splits of GeoQuery split from Shaw et al. [66]. These splits feature structural divergence between training and test examples, e.g., programs have different templates or length. Results in Table 2 (right) shows that grammar prompting can improve upon standard prompting, across all splits (Template, TMCD, Length).

<sup>8</sup>For example, the derivation tree of a subprogram (`attendee_? FindManager(Jean)`) is linearized to `[constraint "(attendee_?" [attendee "FindManager(" [attendee "Jean" "]" ")"])]`, which uses square brackets to encode richer hierarchical information than just the surface form program.

We next assess whether grammar prompting can enable LLMs to make zero-shot use of *unseen functions* (NewFunc) that are not even part of the retrieval set. We set aside 8 functions (smallest, shortest, most, highest, sum, population\_1, count, major) and remove them from the retrieval set, simulating a scenario where new functions are supported in the backend yet no NL-program paired data is available for adapting a semantic parser. Note that for GeoQuery (and Overnight-Blk), we always prepend the full DSL grammar  $G$ —which includes the held-out functions—before the in-context exemplars. Table 2 (right-most column) shows that grammar-prompted LLMs achieve significantly better performance than standard prompting. Our results suggest that the explicit prediction of specialized grammars elicits understanding and reasoning at the grammar level, thereby enabling generalization to unseen functions. We also found that without constrained generation, LLMs were often able to guess functions that did not exist but were nonetheless sensible. An interesting direction is to explore whether LLMs can tackle DSL-open benchmarks such as LARC [1].

**Different base LLMs.** We finally experiment with grammar prompting across different base LLMs. Since GPT-3.5’s 4K token limit is smaller than Codex’s (8K) and GPT-4’s (8K) limits, we use fewer exemplars in these experiments than before (24/8/16 exemplars for GeoQuery/SMCalFlow/Overnight-B respectively). Due to API cost, we limit our experiments to a smaller subset of 100 test examples instead of the full test set.

Table 3 shows that grammar prompting improves upon standard prompting in the majority of the settings. The exceptions are SMCalFlow with GPT-3.5 where both methods performed poorly, and GeoQuery with PaLM 2-L[7], where standard prompting already performed well.

Base LM	Method	GeoQuery	SMCalFlow	Overnight-Blk
Codex	Standard	83	27	63
	Grammar	95	35	66
GPT-3.5	Standard	75	9	49
	Grammar	86	5	67
GPT-4	Standard	85	32	56
	Grammar	98	36	62
PaLM 2-L	Standard	90	14	59
	Grammar	87	17	62

**Table 3:** Results with different base LLMs on a subset of 100 examples sampled from the original test set. GeoQuery and Overnight-Blk show execution accuracy, while SMCalFlow shows program accuracy.

## 4.2 Class-Specific Molecule Generation

We next demonstrate an application of grammar prompting beyond language parsing problems with a molecule generation task. Existing methods for molecule generation typically focus on training specialized neural models using large training sets [45, 37, 15, 2, 79, 61, 19]. We instead follow Guo et al. [29] and explore a few-shot setting where the task is to generate class-specific molecules given a small number of exemplars of that class. Formally, given a small set of molecules  $\{\mathbf{y}_c^{(i)}\}_{i=1}^N$  belonging to a particular molecule class  $c \in \{\text{Acrylates, Chain Extenders, Isocyanates}\}$ , our goal is to generate novel molecules  $\mathbf{y}_c$  of the same class that can be synthesized using existing molecules. Since the in-context examples in this case will only consist of molecules of the same class, the “input”  $x_c^{(i)}$  is the empty string in this case. The data contains 32 Acrylates, 11 Chain Extenders, and 11 Isocyanates (see appendix G of Guo et al. [29]).

While molecules can be more faithfully represented with 3D graph structure, the SMILES string representation [87] remains popular due to its ease of use.<sup>9</sup> The specialized grammars  $G[\mathbf{y}_c]$  (which are specialized from the SMILES grammar) encode various structural properties of the molecule that are specific to the molecule class. Figure 4 shows an example of a specialized grammar and the corresponding molecule in SMILES format. In this example, `ring_closure ::= "1"` specifies the number of rings, and `branch ::= "(" smiles ")"` specifies whether there is a branch.

We test our approach by generating 100 molecules for each class and assessing the quality of the generated molecules. In addition to the standard prompting baseline, we also run the graph grammar baseline from Guo et al. [29] which learns a hypergraph grammar [38] from the given molecules. We use four metrics: *Validity* ( $V$ ), the percentage of chemically valid molecules; *Diversity* ( $D$ ), average pairwise Tanimoto distance over Morgan fingerprints [60]; *Retrosynthesis score* ( $R$ ), the percentage

<sup>9</sup>Note that SMILES does not guarantee that a generated string corresponds to a valid molecule. Using our approach on more advanced string representations such as SELFIES [44] (which guarantee validity) remains an interesting avenue for future work.

Specialized SMILES Grammar for Molecule Generation			
$G[y]$ :	smiles	::=	atom chain branch chain chain   atom chain
	atom	::=	organic_symbol
	organic_symbol	::=	"C"   "N"   "O"
	chain	::=	atom ring_closure bond atom   bond atom   bond atom ring_closure   atom   atom bond atom bond atom   bond atom bond atom
	ring_closure	::=	"1"
	bond	::=	"="
	branch	::=	"(" smiles ")"
<hr/>			
$y$ :	<chem>CC(=C)C(=O)OCCOC1=CC=CC=C1</chem>		

**Figure 4:** Example of a specialized grammar for generating a molecule from the Acrylates class.

Model	Acrylates				Chain Extenders				Isocyanates			
	V	D	R	M	V	D	R	M	V	D	R	M
Graph Grammar [29]	100	0.83	79.0	30.3	100	0.86	72.7	98.3	100	0.93	52.2	82.7
Standard Prompting	87.7	0.73	80.0	76.7	60.3	0.89	72.7	55.7	94.7	0.82	78.0	92.2
Grammar Prompting	98.0	0.74	91.0	93.3	96.3	0.90	86.7	94.0	97.7	0.79	78.0	96.3

**Table 4:** Results for few-shot molecule generation with GPT-3.5. The metrics are validity (V), diversity (D), retrosynthesis score (R) and membership (M). Higher is better for all metrics.

of molecules that are synthesizable from existing molecules, which is computed approximately via the Retro\* model [12]; *Membership (M)*, the percentage of molecules that belong to the desired monomer class. We use GPT-3.5 as the base LLM and sample from the LLM without constrained decoding, as constrained decoding was found to decrease the diversity of samples. See appendix A.2 for the full experimental setup.

**Results.** Table 4 shows that, compared with standard prompting, grammar prompting significantly improves the synthesis of Acrylates and Chain Extenders across all metrics, while yielding mixed results for Isocyanates. Notably, both prompting-based methods outperform the graph grammar baseline in terms of Retro score, possibly due to that LLMs may have been pre-trained on existing datasets of molecules, enabling them to effectively generate synthesizable molecules. In contrast, the baseline method cannot incorporate any external knowledge beyond the 11 or 32 molecules provided. Our preliminary results indicate that LLMs can serve as a useful tool for generating string representations of chemical structures (and potentially other biological/chemical structures).

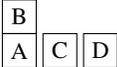
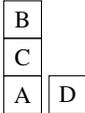
### 4.3 Action Subset Selection for Efficient PDDL Planning

Our final experiments show how grammar prompting can improve the efficiency of classical AI planners. Classical planning is the problem of finding a sequence of actions (i.e., a plan) that goes from an initial state  $s_0$  to a goal state  $s_g$ . An action is represented by a ground operator (e.g., `unstack(block1, block2)` which consists of an operator `unstack` along with two object arguments). We additionally consider *macro-operators* which can potentially speed up planning [9].<sup>10</sup> Planning tasks, along with actions, are represented in Planning Domain Definition Language (PDDL) [27]. We explore how grammar prompted LLMs can help guide classical planning algorithms.

We design specialized grammars to provide guidance to the classical greedy best-first search (GBFS) algorithm [5] by selecting a set of relevant actions. Figure 5 illustrates an example of such a specialized grammar, which captures all the necessary actions for the final plan  $y$  that solves the given task. The process of the guided planning consists of the following steps: (1) given a task, predict a specialized grammar  $G[y]$ ; (2) use the specialized grammar  $G[y]$  to subsequently generate a plan within the restricted action space derived from  $G[y]$ ; (3) initialize GBFS’s priority queue with the LLM-generated plan, and (4) search for the final plan in the restricted action space. Our setup builds upon the idea of using an LLM-generated plan to initialize GBFS from Silver et al. [69], which has a simpler two-step process: (1) given a task, predict a plan via standard prompting, and (2) utilize this plan to guide GBFS. We use their method as our baseline.

Following Silver et al. [69], we create a similar few-shot setting for LLM planning, using 5 tasks as in-context examples and 10 tasks for evaluation from Pyperplan [5]. We test our approach on 3 classic

<sup>10</sup>For example, `pickup-and-stack(A, B)` is a combination of `pickup(A)` and `stack(A, B)`.

Specialized Action Grammar for PDDL Planning			
$s_0$ :		$s_g$ :	
$G[y]$ :	plan	::=	ground_operator+
	ground_operator	::=	"pick_up-and-stack(" block block ")"
			"unstack-and-put_down(" block block ")"
	block	::=	"A"   "B"   "C"
$y$ :	unstack-and-put_down(B, A); pick_up-and-stack(C, A); pick_up-and-stack(B, C)		

**Figure 5:** Example of a specialized grammar for PDDL planning in the Blocks domain. Given an input  $x = (s_0, s_g)$ , the specialized grammar  $G[y]$  only includes necessary actions for solving this task.

Approach	Blocks			Depot			Satellite		
	Created	Expanded	Success	Created	Expanded	Success	Created	Expanded	Success
GBFS + Prim. (No LLM)	360	188	1.0	18047	3870	0.4	8205	150	1.0
Standard + Prim.	348	180	1.0	17597	4039	0.4	6686	78	1.0
Grammar + Prim.	251	124	1.0	15033	3641	0.4	5162	64	1.0
Standard + Macro.	850	16	1.0	1460	56	0.4	4003	27	0.3
Grammar + Macro.	170	9	1.0	2917	127	0.8	3665	46	0.9
Standard + Min Macro.	228	8	1.0	1903	65	0.6	3483	35	0.8

**Table 5:** Results on PDDL planning. Created/Expanded refer to the number of nodes during planning (lower is better). Success refers to success rate (higher is better). Numbers are averaged over three runs using GPT-3.5.

domains in PDDL planning, including Blocks, Depot and Satellite. For the action space, we use either a set of primitive actions (Prim) or an augmented set with macro actions (Macro). In addition to standard prompting, we add two more baselines: (1) No LLM: planning with the entire set of actions; (2) Min Macro: where we construct a minimal set of macro actions for each domain by selecting actions from existing plans for the training tasks. The Min Macro baseline is a domain-specific method to reduce the action space. By comparing to Min Macro, we can verify the effectiveness of instance-specific v.s. domain-specific action selection. See appendix A.3 for more details.

**Results.** We evaluate the efficiency of planning in terms of the number of search nodes created/expanded, as well as the success rate. Table 5 shows the promising performance of LLM-guided planning via grammar prompting. In Blocks, grammar prompting significantly improves efficiency while maintaining 100% success rate. In Depot, grammar prompting with macro actions improves the success rate by 20% over the best competing baseline. In Satellite, using primitive actions yields the best performance with 100% success rate and a reduction of 57% expanded nodes comparing to the No LLM baseline. While our experiments are not intended to compete with the state-of-the-art algorithms for fast planning [20–22, 32, 25, 84], they indicate the promise of LLMs for improving existing planning algorithms.

## 5 Discussion and Limitations

We discuss several limitations of our approach including some negative results. Grammar prompting did not yield any improvements for DSLs that were likely to have been frequently encountered during pretraining (e.g., regular expressions, SQL). Moreover, constrained generation based on specialized grammars led to increased API calls, and was not always beneficial for tasks beyond semantic parsing. For instance, in molecule generation we discovered that enforcing constraints can sometimes result in lower diversity. Additionally, in PDDL planning we observed that the constraints applied to prune objects can sometimes negatively impact performance, suggesting that relevant object selection is still very challenging for LLMs. It may be interesting to explore whether finetuning of moderately-sized LLMs using specialized grammars can lead to better grammar-based models for DSL generation.

On the positive front, our work demonstrates that LLMs have the capacity to understand and generate metalanguages. Working in this “metalanguage space” can be combined with chain-of-thought-style [86] prompts by, for example, manually providing natural language comments to the rules of the specialized grammars. We found this to improve results slightly on semantic parsing (see Figure 6

of appendix A.1). Moreover, many scientific problems can be formally approached by representing hypotheses as DSL programs [71], and DSLs can enable easier encoding of human prior knowledge and scientific principles, providing a foundation for scientific discovery. Recent work shows that state-of-the-art LLMs can follow previously unseen formal systems [75]. Techniques like grammar prompting can widen the scope of scientific problems for which LLMs could be effectively applied by more explicitly accounting for external knowledge and constraints.

## 6 Related Work

**Chain-of-thought prompting.** Grammar prompting extends a recent line of work on improving reasoning capabilities by requesting explicit reasoning steps as part of the prompt [51, 24, 86, 80, 14, 94]. Our approach is closely related to concurrent work on employing symbolic variables as part of the prompt [30, 50, 33, 97, 52], though we are not aware of any existing work that uses formal grammars as the intermediate reasoning step.

**LLMs for program generation and semantic parsing.** Generating programs from natural language specifications, a task often referred to as semantic parsing, is a sub-problem of program synthesis; for surveys, see Kamath and Das [39] and Gulwani et al. [28]. Recent works [8, 89] have explored using LLMs for generating code in general-purpose programming languages (e.g., Python). Our work further extends this line by examining whether LLMs can generate DSL programs, which are intrinsically scarce. There has also been work on using LLMs for tool usage via further training [63] or prompting [56, 77], investigating how model scales [57] and retrievers [96, 46] affect in-context learning for semantic parsing, and constrained decoding [64, 68, 55] for program generation.

**Neural grammars.** Grammar prompting can also be seen as a “fully LLM” instantiation of a line of work on neural parameterizations of symbolic grammars [35, 17, 43, 42, 36, 100, 92, 91, 93]. Indeed, our approach to semantic parsing essentially uses prompt-based learning to define a quasi-synchronous grammar [70, 78] whose rules dynamically depend on the source sentence. Concretely, in contrast to recent works which embed learnable neural components within synchronous grammars [41, 23, 76], grammar prompting relies on the implicit in-context learning capabilities of LLMs for the learning component. (However unlike these works, our conditional grammar does not explicitly align its rules to the subparts of the source sentence).

**Grammar-based molecule generation.** Grammar-based methods have gained significant interest in the realm of molecule generation, offering advantages in interpretability, data-efficiency, and controllability. One line of research involves integrating generic SMILES grammars with neural networks to generate syntactically correct molecules [45, 15]. Another approach centers on data-driven induction of grammars for generation [29, 38]. Our work aligns with the former, viewing grammar prompting as a straightforward method for integrating grammar into an LLM without the need for additional training.

**LLMs for planning.** Recently, LLMs have been increasingly studied in the context of planning for autonomous agents. When given goals expressed in natural language in household environments, earlier works [3, 65, 34, 48] directly prompted LLMs to predict executable actions. However, in PDDL domains, recent works [69, 74] showed that LLMs underperform classical planners if the desired action sequences are very long. Grammar prompting represents a promising strategy for augmenting classical planners with LLMs to get the best of both worlds. Other related efforts include translating between problems and PDDL models [49] and corrective re-prompting [58]. Besides using LLMs, integrating learning and planning has been extensively studied in the past literature, e.g., learning actions [4, 82], skills [84], macro-actions [9], rules [88] and guidance strategies [90, 83, 40] for more efficient planning.

## 7 Conclusion

We propose grammar prompting as a simple approach for improving few-shot DSL generation with large language models. Experiments across a range of structured languages including DSLs for semantic parsing (SMCalFlow, GeoQuery, Overnight), PDDL planning (action DSL), and molecule generation (SMILES), show that grammar prompting can improve upon standard prompting baselines. The encouraging results in semantic parsing indicate its potential to assist LLMs with tool usage, and the promising results in other domains indicate that grammar prompting can enable application of LLMs in domains that intrinsically depend on DSLs.

## Acknowledgments

We thank Jacob Andreas, Gabriel Grand, Linlu Qiu, Tom Silver, and Hunter Lang for helpful discussion and feedback. This study was supported by funds from the Google-MIT research collaborations program and the GIST-MIT joint research program.

## References

- [1] Sam Acquaviva, Yewen Pu, Marta Kryven, Theodoros Sechopoulos, Catherine Wong, Gabrielle Ecanow, Maxwell Nye, Michael Tessler, and Josh Tenenbaum. Communicating natural programs to humans and machines. *Proceedings of NeurIPS*, 35:3731–3743, 2022.
- [2] Walid Ahmad, Elana Simon, Seyone Chithrananda, Gabriel Grand, and Bharath Ramsundar. Chemberta-2: Towards chemical foundation models. *arXiv preprint arXiv:2209.01712*, 2022.
- [3] Michael Ahn, Anthony Brohan, Noah Brown, Yevgen Chebotar, Omar Cortes, Byron David, Chelsea Finn, Keerthana Gopalakrishnan, Karol Hausman, Alex Herzog, et al. Do as I can, not as I say: Grounding language in robotic affordances. *arXiv preprint arXiv:2204.01691*, 2022.
- [4] Diego Aineto, Sergio Jiménez, and Eva Onaindia. Learning STRIPS action models with classical planning. In *Proceedings of the International Conference on Automated Planning and Scheduling*, volume 28, pages 399–407, 2018.
- [5] Yusra Alkhazraji, Matthias Frorath, Markus Grütznert, Malte Helmert, Thomas Liebetraut, Robert Mattmüller, Manuela Ortlieb, Jendrik Seipp, Tobias Springenberg, Philip Stahl, and Jan Wülfing. Pyperplan (v1.3), 2020. URL <https://doi.org/10.5281/zenodo.3701399>.
- [6] Jacob Andreas, John Bufe, David Burkett, Charles Chen, Josh Clausman, Jean Crawford, Kate Crim, Jordan DeLoach, Leah Dorner, Jason Eisner, Hao Fang, Alan Guo, David Hall, Kristin Hayes, Kellie Hill, Diana Ho, Wendy Iwaszuk, Smriti Jha, Dan Klein, Jayant Krishnamurthy, Theo Lanman, Percy Liang, Christopher H. Lin, Ilya Lintsbakh, Andy McGovern, Aleksandr Nisnevich, Adam Pauls, Dmitriy Petters, Brent Read, Dan Roth, Subhro Roy, Jesse Rusak, Beth Short, Div Slomin, Ben Snyder, Stephon Striplin, Yu Su, Zachary Tellman, Sam Thomson, Andrei Vorobev, Izabela Witoszko, Jason Wolfe, Abby Wray, Yuchen Zhang, and Alexander Zotov. Task-oriented dialogue as dataflow synthesis. *Transactions of the Association for Computational Linguistics*, 8:556–571, 2020. doi: 10.1162/tacl\_a\_00333. URL <https://aclanthology.org/2020.tacl-1.36>.
- [7] Rohan Anil, Andrew M Dai, Orhan Firat, Melvin Johnson, Dmitry Lepikhin, Alexandre Passos, Siamak Shakeri, Emanuel Taropa, Paige Bailey, Zhifeng Chen, et al. PaLM 2 technical report. *arXiv preprint arXiv:2305.10403*, 2023.
- [8] Jacob Austin, Augustus Odena, Maxwell Nye, Maarten Bosma, Henryk Michalewski, David Dohan, Ellen Jiang, Carrie Cai, Michael Terry, Quoc Le, et al. Program synthesis with large language models. *arXiv preprint arXiv:2108.07732*, 2021.
- [9] Adi Botea, Markus Enzenberger, Martin Müller, and Jonathan Schaeffer. Macro-ff: Improving AI planning with automatically learned macro-operators. *Journal of Artificial Intelligence Research*, 24:581–621, 2005.
- [10] Tom B Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, et al. Language models are few-shot learners. In *Proceedings of NeurIPS*, 2020.
- [11] Ruisheng Cao, Su Zhu, Chen Liu, Jieyu Li, and Kai Yu. Semantic parsing with dual learning. In *Proceedings of ACL*, pages 51–64, Florence, Italy, July 2019. Association for Computational Linguistics. doi: 10.18653/v1/P19-1007. URL <https://aclanthology.org/P19-1007>.
- [12] Binghong Chen, Chengtao Li, Hanjun Dai, and Le Song. Retro\*: Learning retrosynthetic planning with neural guided A\* search. In *Proceedings of ICML*, pages 1608–1616. PMLR, 2020.

- [13] Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, Alex Ray, Raul Puri, Gretchen Krueger, Michael Petrov, Heidy Khlaaf, Girish Sastry, Pamela Mishkin, Brooke Chan, Scott Gray, Nick Ryder, Mikhail Pavlov, Alethea Power, Lukasz Kaiser, Mohammad Bavarian, Clemens Winter, Philippe Tillet, Felipe Petroski Such, Dave Cummings, Matthias Plappert, Fotios Chantzis, Elizabeth Barnes, Ariel Herbert-Voss, William Hebgen Guss, Alex Nichol, Alex Paino, Nikolas Tezak, Jie Tang, Igor Babuschkin, Suchir Balaji, Shantanu Jain, William Saunders, Christopher Hesse, Andrew N. Carr, Jan Leike, Josh Achiam, Vedant Misra, Evan Morikawa, Alec Radford, Matthew Knight, Miles Brundage, Mira Murati, Katie Mayer, Peter Welinder, Bob McGrew, Dario Amodei, Sam McCandlish, Ilya Sutskever, and Wojciech Zaremba. Evaluating large language models trained on code. 2021.
- [14] Wenhua Chen, Xueguang Ma, Xinyi Wang, and William W Cohen. Program of thoughts prompting: Disentangling computation from reasoning for numerical reasoning tasks. *arXiv preprint arXiv:2211.12588*, 2022.
- [15] Hanjun Dai, Yingtao Tian, Bo Dai, Steven Skiena, and Le Song. Syntax-directed variational autoencoder for structured data. *arXiv preprint arXiv:1802.08786*, 2018.
- [16] David Dohan, Winnie Xu, Aitor Lewkowycz, Jacob Austin, David Bieber, Raphael Gontijo Lopes, Yuhuai Wu, Henryk Michalewski, Rif A. Saurous, Jascha Sohl-dickstein, Kevin Murphy, and Charles Sutton. Language Model Cascades. *arXiv:2207.10342*, 2022.
- [17] Chris Dyer, Adhiguna Kuncoro, Miguel Ballesteros, and Noah A. Smith. Recurrent neural network grammars. In *Proceedings of NAACL*, 2016.
- [18] Jay Earley. An efficient context-free parsing algorithm. *Communications of the ACM*, 13(2): 94–102, 1970.
- [19] Carl Edwards, Tuan Lai, Kevin Ros, Garrett Honke, and Heng Ji. Translation between molecules and natural language. *arXiv preprint arXiv:2204.11817*, 2022.
- [20] Richard E Fikes and Nils J Nilsson. STRIPS: A new approach to the application of theorem proving to problem solving. *Artificial intelligence*, 2(3-4):189–208, 1971.
- [21] Maria Fox and Derek Long. PDDL2. 1: An extension to PDDL for expressing temporal planning domains. *Journal of Artificial Intelligence Research*, 20:61–124, 2003.
- [22] Maria Fox and Derek Long. Modelling mixed discrete-continuous domains for planning. *Journal of Artificial Intelligence Research*, 27:235–297, 2006.
- [23] Dan Friedman, Alexander Wettig, and Danqi Chen. Finding Dataset Shortcuts with Grammar Induction. In *Proceedings of EMNLP*, 2022.
- [24] Luyu Gao, Aman Madaan, Shuyan Zhou, Uri Alon, Pengfei Liu, Yiming Yang, Jamie Callan, and Graham Neubig. PAL: Program-aided Language Models. *arXiv:2211.10435*, 2022.
- [25] Caelan R. Garrett, Tomas Lozano-Perez, and Leslie P. Kaelbling. PDDLStream: Integrating symbolic planners and blackbox samplers. In *International Conference on Automated Planning and Scheduling (ICAPS)*, 2020. URL <https://arxiv.org/abs/1802.08705>.
- [26] Saibo Geng, Martin Josifosky, Maxime Peyrard, and Robert West. Flexible Grammar-Based Constrained Decoding for Language Models. *arXiv:2305.13971*, 2023.
- [27] Malik Ghallab, Adele Howe, Craig Knoblock, Drew McDermott, Ashwin Ram, Manuela Veloso, Daniel Weld, David Wilkins SRI, Anthony Barrett, Dave Christianson, et al. PDDL – the planning domain definition language. *Technical Report CVC TR98003/DCS TR1165. New Haven, CT: Yale Center for Computational Vision and Control.*, 1998.
- [28] Sumit Gulwani, Oleksandr Polozov, Rishabh Singh, et al. Program synthesis. *Foundations and Trends® in Programming Languages*, 4(1-2):1–119, 2017.
- [29] Minghao Guo, Veronika Thost, Beichen Li, Payel Das, Jie Chen, and Wojciech Matusik. Data-efficient graph grammar learning for molecular generation. In *Proceedings of ICLR*, 2022.

- [30] Joy He-Yueya, Gabriel Poesia, Rose E Wang, and Noah D Goodman. Solving math word problems by combining language models with symbolic solvers. *arXiv preprint arXiv:2304.09102*, 2023.
- [31] Jonathan Herzig and Jonathan Berant. Span-based semantic parsing for compositional generalization. In *Proceedings of ACL*, pages 908–921, Online, August 2021. Association for Computational Linguistics. doi: 10.18653/v1/2021.acl-long.74. URL <https://aclanthology.org/2021.acl-long.74>.
- [32] Jörg Hoffmann. The metric-ff planning system: Translating“ignoring delete lists”to numeric state variables. *Journal of Artificial Intelligence Research*, 20:291–341, 2003.
- [33] Hanxu Hu, Hongyuan Lu, Huajian Zhang, Wai Lam, and Yue Zhang. Chain-of-symbol prompting elicits planning in large language models. *arXiv preprint arXiv:2305.10276*, 2023.
- [34] Wenlong Huang, Pieter Abbeel, Deepak Pathak, and Igor Mordatch. Language models as zero-shot planners: Extracting actionable knowledge for embodied agents. In *Proceedings of ICML*, pages 9118–9147. PMLR, 2022.
- [35] Yong Jiang, Wenjuan Han, Kewei Tu, et al. Unsupervised neural dependency parsing. Association for Computational Linguistics (ACL), 2016.
- [36] Lifeng Jin, Finale Doshi-Velez, Timothy Miller, Lane Schwartz, and William Schuler. Unsupervised learning of PCFGs with normalizing flow. In *Proceedings of ACL*, Florence, Italy, July 2019. Association for Computational Linguistics.
- [37] Wengong Jin, Regina Barzilay, and Tommi Jaakkola. Junction tree variational autoencoder for molecular graph generation. In *Proceedings of ICML*, pages 2323–2332. PMLR, 2018.
- [38] Hiroshi Kajino. Molecular hypergraph grammar with its application to molecular optimization. In *Proceedings of ICML*, pages 3183–3191. PMLR, 2019.
- [39] Aishwarya Kamath and Rajarshi Das. A survey on semantic parsing. *arXiv preprint arXiv:1812.00978*, 2018.
- [40] Beomjoon Kim, Zi Wang, Leslie Pack Kaelbling, and Tomás Lozano-Pérez. Learning to guide task and motion planning using score-space representation. *The International Journal of Robotics Research*, 38(7):793–812, 2019.
- [41] Yoon Kim. Sequence-to-sequence learning with latent neural grammars. In *Proceedings of NeurIPS*, pages 26302–26317, 2021.
- [42] Yoon Kim, Chris Dyer, and Alexander Rush. Compound probabilistic context-free grammars for grammar induction. In *Proceedings of ACL*, Florence, Italy, July 2019. Association for Computational Linguistics.
- [43] Yoon Kim, Alexander Rush, Lei Yu, Adhiguna Kuncoro, Chris Dyer, and Gábor Melis. Unsupervised recurrent neural network grammars. In *Proceedings of NAACL*, Minneapolis, Minnesota, June 2019. Association for Computational Linguistics.
- [44] Mario Krenn, Florian Häse, AkshatKumar Nigam, Pascal Friederich, and Alan Aspuru-Guzik. Self-referencing embedded strings (selfies): A 100% robust molecular string representation. *Machine Learning: Science and Technology*, 1(4):045024, oct 2020.
- [45] Matt J Kusner, Brooks Paige, and José Miguel Hernández-Lobato. Grammar variational autoencoder. In *Proceedings of ICML*, pages 1945–1954. PMLR, 2017.
- [46] Itay Levy, Ben Bogin, and Jonathan Berant. Diverse demonstrations improve in-context compositional generalization. *arXiv preprint arXiv:2212.06800*, 2022.
- [47] Yaobo Liang, Chenfei Wu, Ting Song, Wenshan Wu, Yan Xia, Yu Liu, Yang Ou, Shuai Lu, Lei Ji, Shaoguang Mao, Yun Wang, Linjun Shou, Ming Gong, and Nan Duan. TaskMatrix.AI: Completing Tasks by Connecting Foundation Models with Millions of APIs. *arXiv:2303.16434*, 2023.

- [48] Bill Yuchen Lin, Chengsong Huang, Qian Liu, Wenda Gu, Sam Sommerer, and Xiang Ren. On grounded planning for embodied tasks with language models. *arXiv preprint arXiv:2209.00465*, 2022.
- [49] Bo Liu, Yuqian Jiang, Xiaohan Zhang, Qiang Liu, Shiqi Zhang, Joydeep Biswas, and Peter Stone. Llm+ p: Empowering large language models with optimal planning proficiency. *arXiv preprint arXiv:2304.11477*, 2023.
- [50] Qing Lyu, Shreya Havaldar, Adam Stein, Li Zhang, Delip Rao, Eric Wong, Marianna Apidianaki, and Chris Callison-Burch. Faithful chain-of-thought reasoning. *arXiv preprint arXiv:2301.13379*, 2023.
- [51] Maxwell Nye, Anders Andreassen, Guy Gur-Ari, Henryk Witold Michalewski, Jacob Austin, David Bieber, David Martin Dohan, Aitor Lewkowycz, Maarten Paul Bosma, David Luan, Charles Sutton, and Augustus Odena. Show your work: Scratchpads for intermediate computation with language models. *arXiv:2112.00114*, 2021.
- [52] Liangming Pan, Alon Albalak, Xinyi Wang, and William Yang Wang. Logic-LM: Empowering Large Language Models with Symbolic Solvers for Faithful Logical Reasoning. *arXiv:2305.12295*, 2023.
- [53] Bhargavi Paranjape, Scott Lundberg and Sameer Singh, Hannaneh Hajishirzi, Luke Zettlemoyer, and Marco Tulio Ribeiro. ART: Automatic multi-step reasoning and tool-use for large language models. *arXiv:2303.09014*, 2023.
- [54] Emmanouil Antonios Platanios, Adam Pauls, Subhro Roy, Yuchen Zhang, Alexander Kyte, Alan Guo, Sam Thomson, Jayant Krishnamurthy, Jason Wolfe, Jacob Andreas, and Dan Klein. Value-agnostic conversational semantic parsing. In *Proceedings of ACL*, August 2021.
- [55] Gabriel Poesia, Alex Polozov, Vu Le, Ashish Tiwari, Gustavo Soares, Christopher Meek, and Sumit Gulwani. Synchromesh: Reliable Code Generation from Pre-trained Language Models. *Proceedings of ICLR*, 2022.
- [56] Yujia Qin, Shengding Hu, Yankai Lin, Weize Chen, Ning Ding, Ganqu Cui, Zheni Zeng, Yufei Huang, Chaojun Xiao, Chi Han, et al. Tool learning with foundation models. *arXiv preprint arXiv:2304.08354*, 2023.
- [57] Linlu Qiu, Peter Shaw, Panupong Pasupat, Tianze Shi, Jonathan Herzig, Emily Pitler, Fei Sha, and Kristina Toutanova. Evaluating the impact of model scale for compositional generalization in semantic parsing. In *Proceedings of EMNLP*, December 2022.
- [58] Shreyas Sundara Raman, Vanya Cohen, Eric Rosen, Ifrah Idrees, David Paulius, and Stefanie Tellex. Planning with large language models via corrective re-prompting. *arXiv preprint arXiv:2211.09935*, 2022.
- [59] Nils Reimers and Iryna Gurevych. Sentence-bert: Sentence embeddings using siamese bert-networks. *arXiv preprint arXiv:1908.10084*, 2019.
- [60] David Rogers and Mathew Hahn. Extended-connectivity fingerprints. *Journal of chemical information and modeling*, 50(5):742–754, 2010.
- [61] Yu Rong, Yatao Bian, Tingyang Xu, Weiyang Xie, Ying Wei, Wenbing Huang, and Junzhou Huang. Self-supervised graph transformer on large-scale molecular data. *Proceedings of NeurIPS*, 33:12559–12571, 2020.
- [62] Victor Sanh, Albert Webson, Colin Raffel, Stephen H. Bach, et al. Multitask prompted training enables zero-shot task generalization. In *Proceedings of ICLR*, 2022.
- [63] Timo Schick, Jane Dwivedi-Yu, Roberto Dessi, Roberta Raileanu, Maria Lomeli, Luke Zettlemoyer, Nicola Cancedda, and Thomas Scialom. Toolformer: Language Models Can Teach Themselves to Use Tools. *arXiv:2302.04761*, 2023.

- [64] Torsten Scholak, Nathan Schucher, and Dzmitry Bahdanau. PICARD: Parsing incrementally for constrained auto-regressive decoding from language models. In *Proceedings of EMNLP*, Online and Punta Cana, Dominican Republic, November 2021. Association for Computational Linguistics.
- [65] Pratyusha Sharma, Antonio Torralba, and Jacob Andreas. Skill induction and planning with latent language. *arXiv preprint arXiv:2110.01517*, 2021.
- [66] Peter Shaw, Ming-Wei Chang, Panupong Pasupat, and Kristina Toutanova. Compositional generalization and natural language variation: Can a semantic parsing approach handle both? In *Proceedings of ACL*, August 2021.
- [67] Richard Shin and Benjamin Van Durme. Few-shot semantic parsing with language models trained on code. In *Proceedings of NAACL*, pages 5417–5425, Seattle, United States, July 2022. Association for Computational Linguistics. doi: 10.18653/v1/2022.naacl-main.396. URL <https://aclanthology.org/2022.naacl-main.396>.
- [68] Richard Shin, Christopher Lin, Sam Thomson, Charles Chen, Subhro Roy, Emmanouil Antonios Platanios, Adam Pauls, Dan Klein, Jason Eisner, and Benjamin Van Durme. Constrained language models yield few-shot semantic parsers. In *Proceedings of EMNLP*. Association for Computational Linguistics, November 2021.
- [69] Tom Silver, Varun Hariprasad, Reece S Shuttleworth, Nishanth Kumar, Tomás Lozano-Pérez, and Leslie Pack Kaelbling. PDDL planning with pretrained large language models. In *NeurIPS 2022 Foundation Models for Decision Making Workshop*, 2022.
- [70] David Smith and Jason Eisner. Quasi-Synchronous Grammars: Alignment by Soft Projection of Syntactic Dependencies. In *Proceedings on the Workshop on Statistical Machine Translation*, 2006.
- [71] Jennifer J Sun, Megan Tjandrasuwita, Atharva Sehgal, Armando Solar-Lezama, Swarat Chaudhuri, Yisong Yue, and Omar Costilla-Reyes. Neurosymbolic programming for science. *arXiv preprint arXiv:2210.05050*, 2022.
- [72] Didac Suris, Sachit Menon, and Carl Vondrick. ViperGPT: Visual Inference via Python Execution for Reasoning. *arXiv:2303.08128*, 2023.
- [73] Mirac Suzgun, Nathan Scales, Nathanael Scharli, Sebastian Gehrmann, Yi Tay, Hyung Won Chung, Aakanksha Chowdhery, Quoc V. Le, Ed H. Chi, Denny Zhou, and Jason Wei. Challenging big-bench tasks and whether chain-of-thought can solve them. *arXiv:2210.09261*, 2022.
- [74] Karthik Valmeekam, Alberto Olmo, Sarath Sreedharan, and Subbarao Kambhampati. Large language models still can’t plan (a benchmark for llms on planning and reasoning about change). *arXiv preprint arXiv:2206.10498*, 2022.
- [75] Gregor vom Scheidt. Experimental results from applying GPT-4 to an unpublished formal language. *arXiv:2305.12196*, 2023.
- [76] Bailin Wang, Ivan Titov, Jacob Andreas, and Yoon Kim. Hierarchical Phrase-based Sequence-to-Sequence Learning. In *Proceedings of EMNLP*, 2022.
- [77] Guanzhi Wang, Yuqi Xie, Yunfan Jiang, Ajay Mandlekar, Chaowei Xiao, Yuke Zhu, Linxi Fan, and Anima Anandkumar. Voyager: An open-ended embodied agent with large language models. *arXiv preprint arXiv:2305.16291*, 2023.
- [78] Mengqiu Wang, Noah A. Smith, and Teruko Mitamura. What is the Jeopardy Model? A Quasi-Synchronous Grammar for QA. In *Proceedings of EMNLP*, 2007.
- [79] Sheng Wang, Yuzhi Guo, Yuhong Wang, Hongmao Sun, and Junzhou Huang. Smiles-bert: large scale unsupervised pre-training for molecular property prediction. In *Proceedings of the 10th ACM international conference on bioinformatics, computational biology and health informatics*, pages 429–436, 2019.

- [80] Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, Sharan Narang, Aakanksha Chowdhery, and Denny Zhou. Self-consistency improves chain of thought reasoning in language models. In *Proceedings of ICLR*, 2023.
- [81] Yushi Wang, Jonathan Berant, and Percy Liang. Building a semantic parser overnight. In *Proceedings of ACL*, Beijing, China, July 2015.
- [82] Zi Wang, Stefanie Jegelka, Leslie Pack Kaelbling, and Tomás Lozano-Pérez. Focused model-learning and planning for non-Gaussian continuous state-action systems. In *2017 IEEE International conference on robotics and automation (ICRA)*, pages 3754–3761. IEEE, 2017.
- [83] Zi Wang, Caelan Reed Garrett, Leslie Pack Kaelbling, and Tomás Lozano-Pérez. Active model learning and diverse action sampling for task and motion planning. In *2018 IEEE/RSJ International Conference on Intelligent Robots and Systems (IROS)*, pages 4107–4114. IEEE, 2018.
- [84] Zi Wang, Caelan Reed Garrett, Leslie Pack Kaelbling, and Tomás Lozano-Pérez. Learning compositional models of robot skills for task and motion planning. *The International Journal of Robotics Research*, 40(6-7):866–894, 2021.
- [85] Jason Wei, Maarten Paul Bosma, Vincent Zhao, Kelvin Guu, Adams Wei Yu, Brian Lester, Nan Du, Andrew Mingbo Dai, and Quoc V. Le. Finetuned language models are zero-shot learners. In *Proceedings of ICLR*, 2022.
- [86] Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, Brian Ichter, Fei Xia, Ed Chi, Quoc Le, and Denny Zhou. Chain of thought prompting elicits reasoning in large language models. In *Proceedings of NeurIPS*, 2022.
- [87] David Weininger. Smiles, a chemical language and information system. 1. introduction to methodology and encoding rules. *Journal of chemical information and computer sciences*, 28(1):31–36, 1988.
- [88] Victoria Xia, Zi Wang, and Leslie Pack Kaelbling. Learning sparse relational transition models. In *International Conference on Learning Representations*, 2019.
- [89] Frank F Xu, Uri Alon, Graham Neubig, and Vincent Josua Hellendoorn. A systematic evaluation of large language models of code. In *Proceedings of the 6th ACM SIGPLAN International Symposium on Machine Programming*, pages 1–10, 2022.
- [90] Ryan Yang, Tom Silver, Aidan Curtis, Tomas Lozano-Perez, and Leslie Pack Kaelbling. Pg3: Policy-guided planning for generalized policy generation. *arXiv preprint arXiv:2204.10420*, 2022.
- [91] Songlin Yang, Yanpeng Zhao, and Kewei Tu. Neural bi-lexicalized PCFG induction. In *Proceedings of ACL*, Online, August 2021. Association for Computational Linguistics.
- [92] Songlin Yang, Yanpeng Zhao, and Kewei Tu. PCFGs can do better: Inducing probabilistic context-free grammars with many symbols. In *Proceedings of NAACL*, 2021.
- [93] Songlin Yang, Roger P Levy, and Yoon Kim. Unsupervised discontinuous constituency parsing with mildly context-sensitive grammars. *arXiv preprint arXiv:2212.09140*, 2022.
- [94] Shunyu Yao, Dian Yu, Jeffrey Zhao, Izhak Shafran, Thomas L Griffiths, Yuan Cao, and Karthik Narasimhan. Tree of thoughts: Deliberate problem solving with large language models. *arXiv preprint arXiv:2305.10601*, 2023.
- [95] Jiacheng Ye, Zhiyong Wu, Jiangtao Feng, Tao Yu, and Lingpeng Kong. Compositional exemplars for in-context learning. *arXiv preprint arXiv:2302.05698*, 2023.
- [96] Xi Ye, Qiaochu Chen, Isil Dillig, and Greg Durrett. Benchmarking multimodal regex synthesis with complex structures. In *Proceedings of ACL*. Association for Computational Linguistics, July 2020.
- [97] Xi Ye, Qiaochu Chen, Isil Dillig, and Greg Durrett. Satisfiability-aided language models using declarative prompting. *arXiv preprint arXiv:2305.09656*, 2023.

- [98] Pengcheng Yin, Hao Fang, Graham Neubig, Adam Pauls, Emmanouil Antonios Platanios, Yu Su, Sam Thomson, and Jacob Andreas. Compositional generalization for neural semantic parsing via span-level supervised attention. *Proceedings of ACL*, 2021.
- [99] John M Zelle and Raymond J Mooney. Learning to parse database queries using inductive logic programming. In *Proceedings of the national conference on artificial intelligence*, pages 1050–1055, 1996.
- [100] Hao Zhu, Yonatan Bisk, and Graham Neubig. The return of lexical dependencies: Neural lexicalized PCFGs. *Transactions of the Association for Computational Linguistics*, 8, 2020.

	Few-Shot			Retrieval-based			GeoQuery Out-of-Dist.			
	GeoQuery	SMCalflow	Overnight-Blk	GeoQuery	SMCalflow	Overnight-Blk	Template	TMCD	Length	NewFunc
Train	32	16	32	560	128	1436	441	440	440	453
Test	280	360	399	280	360	399	439	440	440	447

**Table 6:** Statistics of the splits used for experiments on semantic parsing.

Approach	GeoQuery	SMCalFlow	Overnight-B
Standard Prompting (unconstrained decoding)	81.5 (1.0)	46.4 (1.0)	54.7 (1.0)
<i>w. constrained decoding (<math>\hat{\mathbf{y}} \in L(G)</math>)</i>	81.8 (4.3)	49.2 (5.6)	54.7 (1.6)
Linearized Derivation Tree Prompting	77.5 (1.0)	50.0 (1.0)	56.4 (1.0)
Grammar Prompting (unconstrained decoding)	87.5 (1.0)	50.8 (1.0)	57.4 (1.0)
<i>w. grammar constraint (<math>\hat{G} \subseteq G</math>)</i>	88.6 (3.0)	51.3 (3.0)	60.4 (1.4)
<i>w. grammar and program constraint (<math>\hat{\mathbf{y}} \in L(\hat{G})</math>)</i>	88.9 (3.3)	52.4 (3.3)	60.9 (2.8)
<i>w. oracle grammar (<math>\hat{G} = G[\mathbf{y}]</math>)</i>	96.1 (1.3)	80.0 (1.0)	94.2 (1.0)
<i>w. oracle grammar + program constraint</i>	96.8 (2.1)	83.6 (2.6)	96.5 (1.0)

**Table 7:** Extended results which show the number of Codex calls per example on few-shot semantic parsing in brackets. Columns in grey show program accuracy, while white columns others indicate execution accuracy.

## A Experiment Details

### A.1 Semantic Parsing

**Statistics and Splits.** We show the statistics for the splits used for the experiments in Table 6.

For GeoQuery, we utilize the standard split from Zelle and Mooney [99] in the retrieval-based setting and the Template, TMCD, and Length splits from Shaw et al. [66]. We randomly sample 32 examples from the training set of the standard split to create the few-shot split. To generate the NewFunc split, we designate examples utilizing the following eight functions as test examples: smallest, shortest, most, highest, sum, population\_1, count, major; the remaining examples are incorporated into the training set.

For SMCaFlow, we adopt the 16-shot and 128-shot cross-domain settings from Yin et al. [98] as our few-shot and retrieval-based settings, respectively. It is noteworthy that the training set of the original splits contains approximately 25k in-domain training examples. Previous work [96, 57] utilized these examples as their retrieval set. However, in our experiments, we found that incorporating in-domain examples did not enhance performance. Consequently, we use 16/128 cross-domain examples as our training set in the few-shot and retrieval settings, respectively. For all experiments on SMCaFlow, we used the preprocessed version from Qiu et al. [57], which employs a more concise LISPRESS format [54] than the original version [98]. This format aligns with Ye et al. [96] for a fair comparison.

For Overnight-Blocks, we employ the standard split from Wang et al. [81] in the retrieval setting. We randomly sample 32 examples from the training set of the standard split to create the few-shot split.

**Scoring Functions for Constrained Generation.** For each candidate continuation  $w \in \Sigma[\mathbf{y}_{\text{prefix}}]$  for correction, we first form a partial program via concatenation  $\mathbf{y}_{\text{prefix}} \cdot w$  and then feed it into Codex to obtain the score for the candidate via

$$\log P_{\text{LLM}}(w \mid \hat{G}, \mathbf{x}, \mathbf{y}_{\text{prefix}}, (\mathbf{x}^{(i)}, G[\mathbf{y}^{(i)}], \mathbf{y}^{(i)})_{i=1}^N).$$

In the case where  $w$  consists of multiple BPE tokens, e.g., FindManger( is tokenized into Find, Manager, and (, we average the token-level log-likelihood to obtain a candidate-level score. However, when the size of  $\Sigma[\mathbf{y}_{\text{prefix}}]$  exceeds 16, invoking Codex for each candidate becomes too expensive. To address this issue, we employ SentenceBERT to select 16 most plausible candidates first via a dot product,

$$(\text{SentenceBERT}(\hat{\mathbf{y}}_t))^\top (\text{SentenceBERT}(\mathbf{y}_{\text{prefix}} \cdot w)),$$

where SentenceBERT yields the embeddings for the incorrect prediction  $\hat{\mathbf{y}}_t$  and a candidate of corrected partial program  $\mathbf{y}_{\text{prefix}} \cdot w$ .

Molecule Class	Temperature	Presence Penalty	Frequency Penalty
<i>Sampling specialized grammars <math>\hat{G}</math></i>			
Acrylates	0.6	0.1	0.1
Chain Extenders	0.6	0.1	0.1
Isocyanates	0.6	0.4	0.4
<i>Sampling molecules <math>\hat{y}</math></i>			
Acrylates	0.6	0.1	0.1
Chain Extenders	0.6	0.1	0.1
Isocyanates	0.3	0.4	0.4

**Table 8:** Hyperparameters for sampling specialized grammars  $\hat{G}$  (top) and the molecules  $\hat{y}$  in grammar prompting for molecule generation. Standard prompting uses the same hyperparameters for  $y$ .

The functionality of obtaining the log-likelihood for a candidate continuation given a prefix is applicable via Codex APIs <sup>11</sup> via setting `logprobs=True` and `echo=True`. Unfortunately, subsequent models (e.g., GPT-3.5 and GPT-4) disable such functionality. As a workaround, we simply use the scoring function based on SentenceBERT to directly select the best candidate in our PDDL planning experiments.

**Cost Efficiency.** We assess various decoding strategies for their cost efficiency, focusing on the number of API calls. The number of Codex calls resulting from the few-shot semantic parsing experiments is presented in Figure 7, alongside the corresponding accuracy metrics. The results indicate that standard prompting under constrained decoding leads to a significantly higher number of Codex calls. Similarly, grammar-prompting with constraints also results in an increased number of Codex calls. However, when employing both grammar and program constraints, the number of calls decreases meaningfully in comparison to standard prompting under constrained decoding. Future work might consider exploring strategies for more cost-efficient constrained decoding.

**Grammar Prompting with Annotated Rules.** We have additionally experimented with enhancing BNF rules by appending natural language comments. As illustrated in Figure 6, we pair each BNF rule with its corresponding natural language phrases extracted from the given query  $x$ . These manually annotated comments yield an explicit correspondence between natural language phrases and their corresponding BNF rules, thereby better facilitating interpretation and application of the grammars for generating programs. When employing the augmented grammar prompting, we noticed marginal improvements on SMCaFlow (+1.0%) and Overnight-Blocks (0.5%), with no observed enhancement on GeoQuery. While the gains might not appear significant, this predicted alignment could potentially contribute to improved interpretability and further constraints on generation. For instance, the phrase “someone’s manager” should consistently trigger the function `FindManager()`. We leave the exploration of utilizing the augmented rules for future work.

## A.2 Molecule Generation

**Sampling Procedure** Different from semantic parsing and PDDL planning, where the most probable program  $y$  needs to be found via `arg max` inference, molecule generation has empty specification  $x$  and requires sampling from a prompting-based distribution. The sampling procedure for grammar prompting consists of three stages: (1) we randomly sample a permutation of given molecules, denoted as  $\pi$ , (2) based on a prompt formed by the permutation, we sample a specialized grammar  $\hat{G}$  via

$$\hat{G} \sim P_{\text{LLM}}(G' \mid x, (G[y^{(\pi[i])}], y^{(\pi[i])})_{i=1}^N),$$

iii) we finally obtain the molecule conditioned  $\hat{G}$ ,

$$\hat{y} \sim P_{\text{LLM}}(y \mid \hat{G}, (G[y^{(\pi[i])}], y^{(\pi[i])})_{i=1}^N).$$

We list the hyperparameters used for the sampling procedure in for (2) in Table 8 (top) and for (3) in Table 8 (bottom).

<sup>11</sup><https://learn.microsoft.com/en-us/azure/cognitive-services/openai/reference>

LLM Prompt			
You are an expert programmer, and you need to write a program for the given natural language query. First, you should write a grammar that contains all the necessary BNF rules. Then, you should write programs that conform to your predicted rules.			
$x^{(1)}$ : find the meeting on Wednesday with Bob and Carol			
$G[y^{(1)}]$ :	event	::= "QueryEvent(" constraint ")"	<i>find the meeting</i>
	constraint	::= "&" constraint constraint "	
		"(start_?" date ")"	<i>on ...</i>
		"(attendee_?" attendee attendee ")"	<i>with ...</i>
	date	::= "Wednesday"	<i>Wednesday</i>
	attendee	::= "Bob"   "Carol"	<i>Bob and Carol</i>
$y^{(1)}$ : QueryEvent(& (start_? Wednesday)(attendee_? Bob Carol))			
$x$ : Add meeting with Jean's manager on Wednesday at 3PM			
LLM Output			
$\hat{G}$ :	event	::= "CreateEvent(" constraint ")"	<i>Add meeting</i>
	constraint	::= "&" constraint constraint "	
		"(start_?" date time ")"	<i>on ... at ...</i>
		"(attendee_?" attendee ")"	<i>with ...</i>
	date	::= "Wednesday"	<i>Wednesday</i>
	time	::= "NumberPM(3)"	<i>3PM</i>
	attendee	::= "FindManager(" attendee ")   "Jean"	<i>Jean's manager</i>
$\hat{y}$ : CreateEvent(& (start_? Wednesday NumberPM(3))(attendee_? FindManager(Jean)))			

**Figure 6:** Example of grammar prompting where BNF grammars are additionally annotated with natural language comments (shown in green). These manually curated comments provide a detailed mapping between natural language phrases and their corresponding BNF rules, thereby better facilitating interpretation and application of the grammars for generating programs. We manually craft and add these comments to the few-shot prompts (top). The model predicts this during inference (bottom).

In comparison, the sampling procedure for standard prompting only consists of two stages: (1) we randomly sample a permutation of given molecules, denoted as  $\pi$ , (2) based on a prompt formed by the permutation, we directly sample a molecule via

$$\hat{y} \sim P_{\text{LLM}}(y | (y^{(\pi[i])})_{i=1}^N).$$

The hyperparameters used for Step (2) is the same as in grammar prompting and shown in Table 8 (bottom).

While we observed that Earley-based constrained generation enhances grammar prompting in terms of improving validity, other metrics, such as the retrosynthesis score, decreased significantly. This discrepancy could be attributed to the fact that existing LLMs, due to their limited exposure to molecules represented in SMILES format, struggle with comprehending and applying the BNF grammar rules of SMILES. Overall, our current findings serve as preliminary evidence that grammar prompting can tap into the capacity of LLMs to understand and apply BNF rules. However such capacity still remains limited in text-focused LLMs.

### A.3 PDDL Planning

**Restricted Action Space** The specialized grammar defined for PDDL planning essentially delineates a constrained action space that includes necessary actions and their associated objects. Our empirical results found that limiting the classical GBFS planner to the objects selected by a specialized grammar proved too restrictive, yielding beneficial results only within the Blocks domain. Therefore, we decided to remove this limitation, thus expanding the action space of GBFS to contain the actions predicted from the grammar with an unrestricted range of objects.

## B Prompts

Figures 7, 8, and 9 demonstrate the prompts with grammars, based on actual examples in the SMCaFlow, GeoQuery, and Overnight datasets respectively. Because the general grammar of SMCaFlow is long (around 4k tokens), we do not include it within the prompt. For GeoQuery and Overnight, the general grammar is integrated as part of the instruction within the prompt. In the context of molecule generation, the general grammar for SMILES<sup>12</sup> is also included. Figures 10 and 11 demonstrate the prompts with action DSLs for PDDL planning.

---

<sup>12</sup><https://metamolecular.com/cheminformatics/smiles/formal-grammar/>

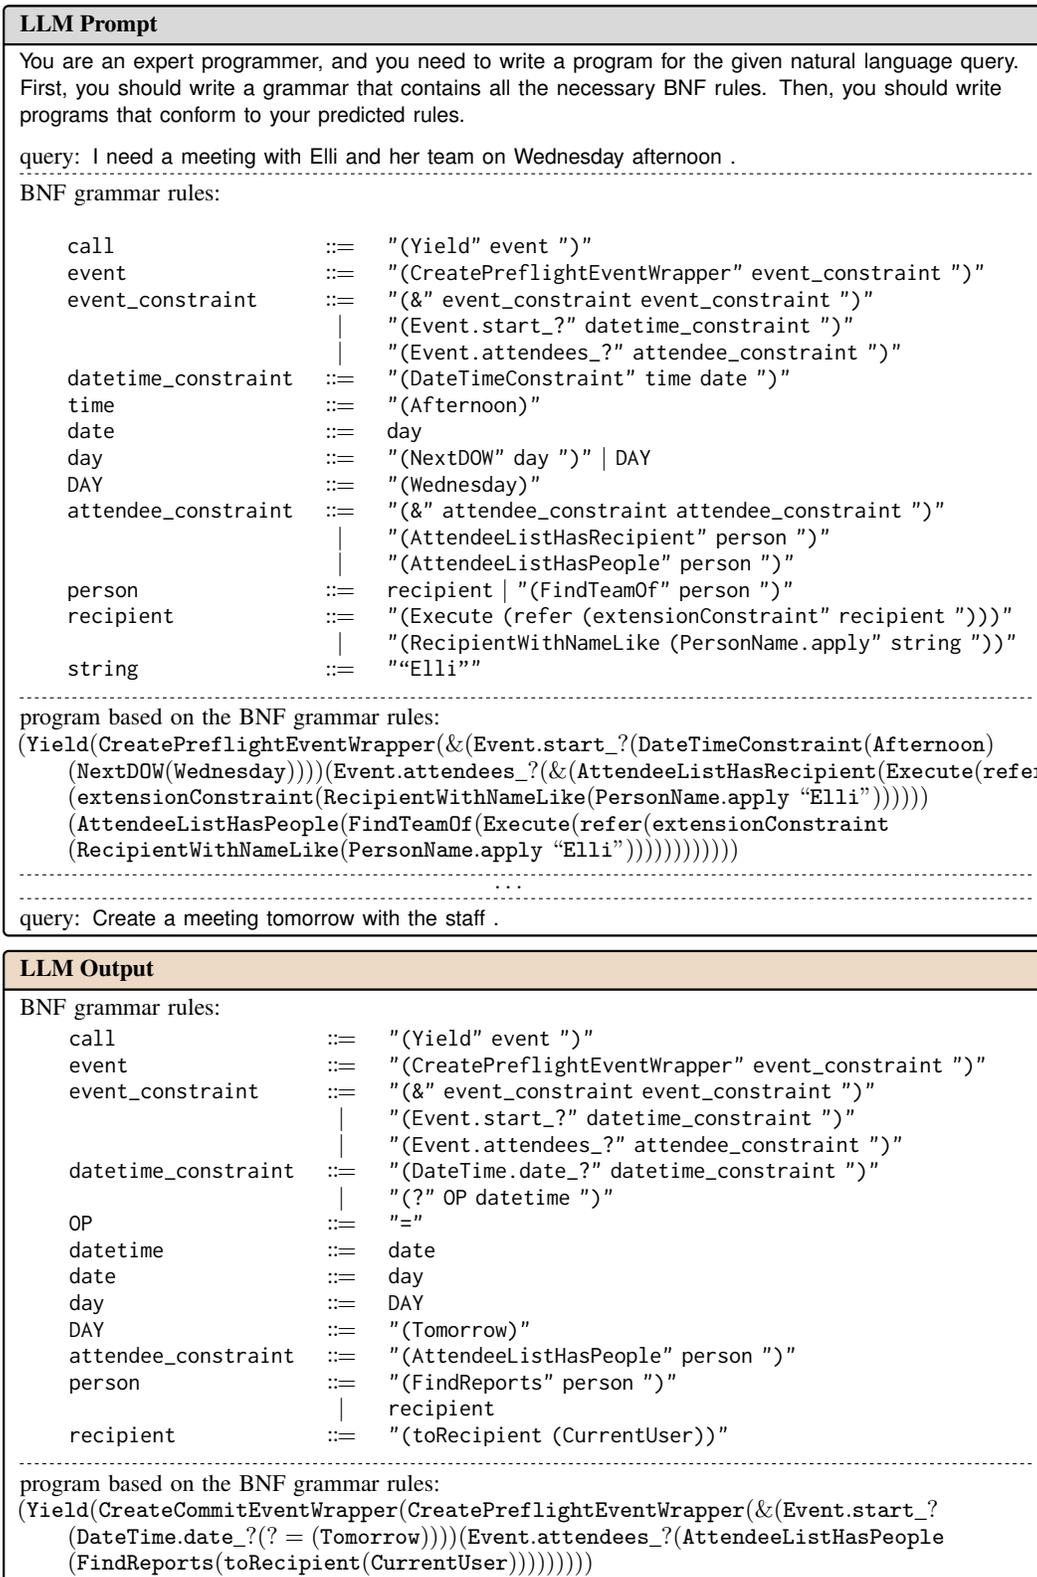


Figure 7: Prompt with real examples from the SMCaFlow dataset.

**LLM Prompt**

You are an expert programmer, and you need to write a program for the given natural language query. First, you should write grammar rules by choosing from the following BNF rules. Then, you should write programs that conform to your predicted rules.

---

```
[BEGIN RULES]
    query      ::= "answer(" answer_type ")"
    answer_type ::= city | state | num | place | river | country
    city       ::= "city(" city ")"
                | "cityid(" CITYNAME ", " STATEABBREV "')"
                | "cityid(" CITYNAME ", _)"
                | "capital(" city ")"
                | "major(" city ")"
                | "capital_1(" state ")"
                | "loc_2(" state ")"
                | "loc_2(" country ")"
                | "largest(" city ")"
                | "smallest(" city ")"
                | "intersection(" city ", " city ")"
                | "exclude(" city ", " city ")"
                | "largest_one(population_1(" city "))"
                | "largest_one(density_1(" city "))"
                | "smallest_one(population_1(" city "))"
                | "smallest_one(density_1(" city "))"
                | ALL_CITY
                ...
[END RULES]
```

---

query: what states border hawaii ?  
BNF grammar rules:

```

    query      ::= "answer(" answer_type ")"
    answer_type ::= state
    state      ::= "state(" state ")"
                | "next_to_2(" state ")"
                | "stateid(" STATENAME "')"
    STATENAME  ::= "hawaii"
```

---

program based on the BNF grammar rules:  
answer(state(next\_to\_2(stateid('hawaii'))))

---

query: how many major cities are in arizona ?

**LLM Output**

BNF grammar rules:

```

    query      ::= "answer(" answer_type ")"
    answer_type ::= num
    num        ::= "count(" city ")"
    city       ::= "major(" city ")"
                | "city(" city ")"
                | "loc_2(" state ")"
    state      ::= "stateid(" STATENAME "')"
    STATENAME  ::= "arizona"
```

---

program based on the BNF grammar rules:  
answer(count(major(city(loc\_2(stateid('arizona'))))))

**Figure 8:** Prompt with real examples from the GeoQuery dataset.

**LLM Prompt**

You are an expert programmer, and you need to write a program for the given natural language query. First, you should write grammar rules by choosing from the following BNF rules. Then, you should write programs that conform to your predicted rules.

---

[BEGIN RULES]

```

list_value ::= "(listValue" list_value ")"
           |(filter" list_value PROPERTY ")
           |(filter" list_value PROPERTY OP list_value ")
           |(superlative" list_value AGGREGATE "
           (ensureNumericProperty" PROPERTY ")
           |(countSuperlative" list_value AGGREGATE PROPERTY ")
           |("_size" list_value ")
           |(aggregate" AGGREGATE list_value ")
           |(getProperty" list_value PROPERTY ")
           |(getProperty (singleton" SINGLETON_VALUE ") !type)"
           |(concat" ENTITY_VALUE ENTITY_VALUE ")
           |(concat" NUMBER_VALUE NUMBER_VALUE ")
           |ENTITY_VALUE | NUMBER_VALUE

PROPERTY ::= "shape" | "color" | "length" | "is_special" | "width"
           | "height" | "left" | "right" | "above" | "below"
           |(reverse left)" | "(reverse right)"
           |(reverse above)" | "(reverse below)"

SINGLETON_VALUE ::= "en.block" | "en.shape" | "en.color"
ENTITY_VALUE ::= "en.block.block1" | "en.block.block2" | "en.shape.pyramid"
              | "en.shape.cube" | "en.color.red" | "en.color.green"
NUMBER_VALUE ::= "3 en.inch" | "6 en.inch" | "2"
OP ::= "=" | ">" | "<" | ">=" | "<=" | "!="
AGGREGATE ::= "sum" | "max" | "min" | "avg"

```

[END RULES]

query: which brick is taller than brick 1

BNF grammar rules:

```

list_value ::= "(listValue" list_value ")
           |(aggregate" AGGREGATE list_value ")
           |(getProperty" list_value PROPERTY ")
           |(getProperty (singleton" SINGLETON_VALUE ") !type)"

AGGREGATE ::= "avg"
SINGLETON_VALUE ::= "en.block"
PROPERTY ::= "width"

```

---

program based on the BNF grammar rules:  
(listValue(aggregate avg (getProperty(getProperty(singleton en.block) !type) width)))

---

query: look right of block 1 and get the special block

**LLM Output**

BNF grammar rules:

```

list_value ::= "(listValue" list_value ")
           |(filter" list_value PROPERTY OP list_value ")
           |(filter" list_value PROPERTY ")
           |(getProperty (singleton" SINGLETON_VALUE ") !type)"
           |ENTITY_VALUE

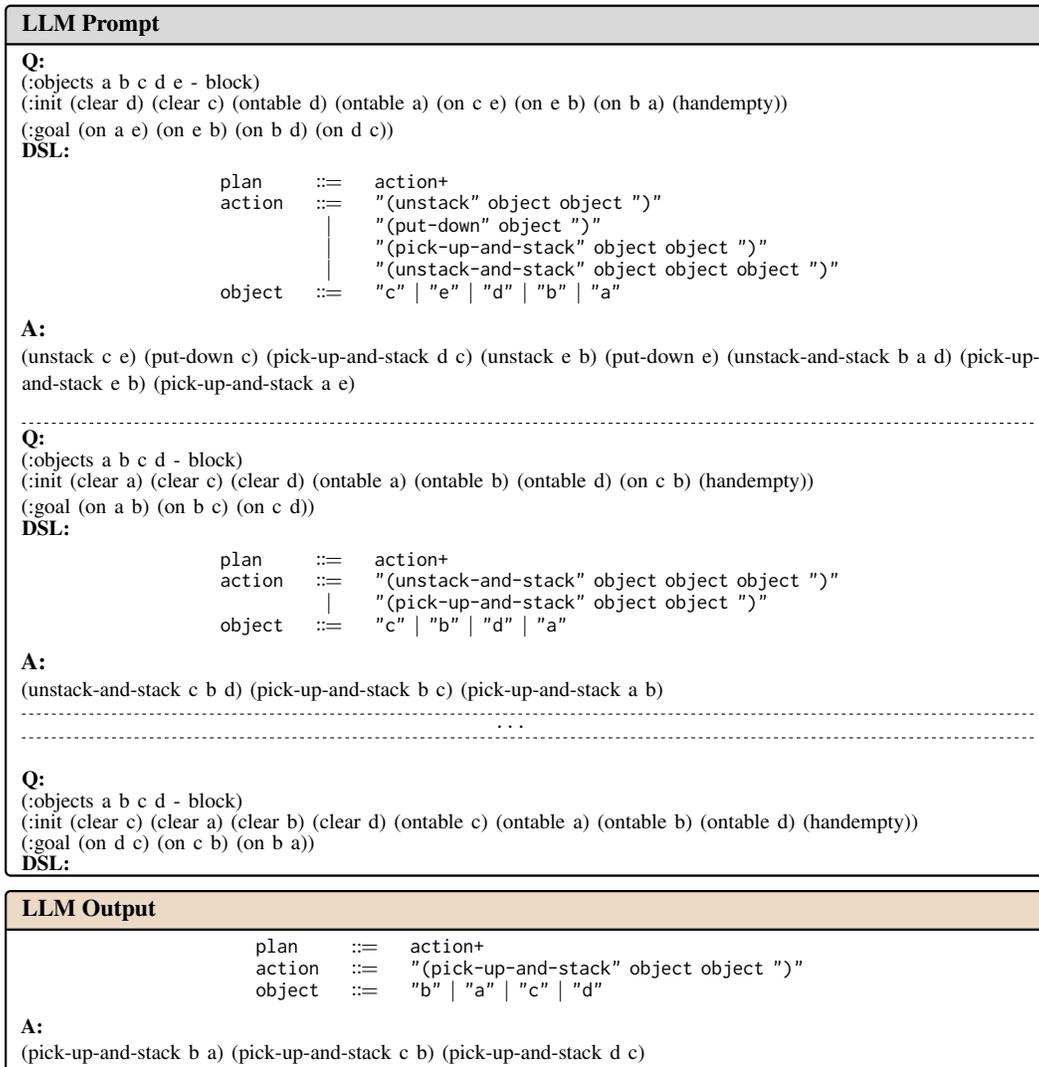
SINGLETON_VALUE ::= "en.block"
PROPERTY ::= "is_special"
           |(reverse left)"
OP ::= "="
ENTITY_VALUE ::= "en.block.block1"

```

---

program based on the BNF grammar rules:  
(listValue(filter(filter(getProperty(singleton en.block) !type) is\_special)
(reverse left) = en.block.block1))

Figure 9: Prompt with real examples from the Overnight dataset. Some long rules are omitted to fit the space.



**Figure 10:** Prompt with real examples in the Blocks domain from Pyperplan. The prompt template follows [69].



Figure 11: Prompt with real examples in the Depot domain from Pyperplan. The prompt template follows [69].