

Distilling Script Knowledge from Large Language Models for Constrained Language Planning

Siyu Yuan[♡], Jiangjie Chen^{♣*}, Ziquan Fu^{♣†}, Xuyang Ge[♣],
Soham Shah[◇], Charles Robert Jankowski[◇], Yanghua Xiao^{♣¶}, Deqing Yang^{♡*}

[♡]School of Data Science, Fudan University

[♣]Shanghai Key Laboratory of Data Science, School of Computer Science, Fudan University

[♣]System Inc. [◇]Brain Technologies, Inc.

[¶]Fudan-Aishu Cognitive Intelligence Joint Research Center

siyuan21@m.fudan.edu.cn, jjchen19@fudan.edu.cn

frank@system.com, {sshah, cjankowski}@brain.im

{xyge20, shawyh, yangdeqing}@fudan.edu.cn

Abstract

In everyday life, humans often plan their actions by following step-by-step instructions in the form of goal-oriented scripts. Previous work has exploited language models (LMs) to plan for abstract goals of stereotypical activities (e.g., “make a cake”), but leaves more specific goals with multi-facet constraints understudied (e.g., “make a cake for diabetics”). In this paper, we define the task of constrained language planning for the first time. We propose an over-generate-then-filter approach to improve large language models (LLMs) on this task, and use it to distill a novel constrained language planning dataset, CoScript, which consists of 55,000 scripts. Empirical results demonstrate that our method significantly improves the constrained language planning ability of LLMs, especially on constraint faithfulness. Furthermore, CoScript is demonstrated to be quite effective in endowing smaller LMs with constrained language planning ability.¹

1 Introduction

To accomplish everyday goals, humans usually plan their actions in accordance with step-by-step instructions. Such instructions are discovered as *goal-oriented scripts* (Schank and Abelson, 1975, 2013), involving a set of prototypical event sequences to achieve goals. For the example in Figure 1, to achieve the goal (*make a cake*), one usually has to follow certain steps of instructions, e.g., *gather ingredients*, *preheat the oven*, etc. The planning for such step-by-step scripts chains up reasoning toward complex goals (Abelson, 1976; Wei et al., 2022). Therefore, the automation of planning envisions more intelligent and reasonable AI

*Corresponding authors.

†Work done while at Brain Technologies, Inc.

¹Resources of this paper can be found at <https://github.com/siyuyuan/coscript>.

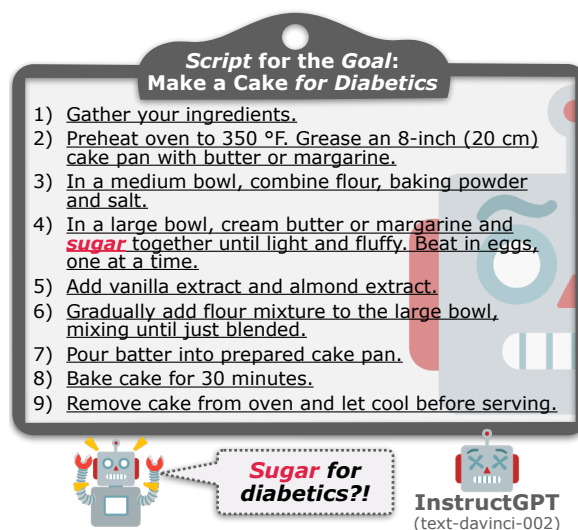


Figure 1: A list of steps InstructGPT generates to plan for the goal “make a cake for diabetics”. InstructGPT mistakenly adds sugar to the cake, which is unfit for diabetic patients. This example shows that InstructGPT sometimes cannot effectively and faithfully script for a *specific* goal with fine-grained constraints.

systems in various domains, such as executable robotic systems (Kovalchuk et al., 2021; Huang et al., 2022) and reasoning systems for problem-solving (Wei et al., 2022; Wang et al., 2022).

Recent studies have identified that language models (LMs) can be used to plan scripts (Sancheti and Rudinger, 2022). Previous work (Huang et al., 2022) has shown that large language models (LLMs), such as GPT-3 (Brown et al., 2020) InstructGPT (Ouyang et al., 2022) and PaLM (Chowdhery et al., 2022), can effectively decompose goals into procedural steps in a zero-/few-shot manner. To train specialized models, researchers have proposed datasets for the automatic understanding and generation of script knowledge (Schank and Abelson, 1975; Regneri et al., 2010; Wanzare et al., 2016; Lyu et al., 2021; Sak-

aguchi et al., 2021). However, previous work mainly focuses on planning for the abstract goals of stereotypical activities (Abelson, 1976). Planning for goals with specific constraints (e.g., for diabetics) still remains under-studied.

In this paper, we define the problem of *constrained language planning*, which imposes different constraints on the goals of planning. An *abstract goal*, for example, *make a cake*, can be inherited by different real-life *specific goals* with multi-faceted *constraints*. A cake can be made for 1) different ingredients (e.g., *chocolate* or *vanilla*); 2) various tools (e.g., with a *microwave* or an *oven*); or 3) different purposes (e.g., for a *wedding* or a *birthday party*). A good planner should write scripts that are reasonable and faithful to constraints. However, LLMs sometimes do not plan faithfully toward the constraints. As showcased in Figure 1, InstructGPT suggests adding sugar to the cake for diabetic patients. Also, due to a shortage of datasets for constrained language planning, the ability of smaller but specialized models to plan with specific constraints has been underexplored.

In this paper, we aim to evaluate and improve the constrained language planning ability of LLMs, while distilling a dataset from LLMs to train specialized models. Our empirical study finds that LLMs tend to plan fluently but unfaithfully to the constraints. Thus, we employ an over-generate-then-filter approach (Wiegreffe et al., 2022) to satisfy the quality of the generated scripts to constraints. The main idea is to select high-quality ones from multiple generated scripts. Then, we use LLMs (e.g., InstructGPT) with this approach to generate a dataset for constrained language planning, which inherits the idea of symbolic knowledge distillation from models (West et al., 2022). We thus arrive at a **Constrained Script** dataset, i.e., *CoScript*, which consists of 55,000 high-quality scripts with specific goals and steps. Experiments show that, when trained on *CoScript*, smaller models such as T5 (Raffel et al., 2020) can achieve good performance, even surpassing that of LLMs.

Our contributions are summarized as follows: 1) To our knowledge, we are the first to establish the constrained language planning problem, which advances language planning toward more specific goals. 2) We evaluate the few-shot constrained language planning ability of LLMs and develop an over-generate-then-filter method for LLMs, resulting in a 26% increase in accuracy. 3) Based

on our method, we use LLMs to generate a high-quality script dataset (*CoScript*) for constrained language planning. By leveraging the *CoScript*, we endow specialized and smaller models with constrained language planning ability, which achieves comparable performance to that of LLMs.

2 Related Work

Language Planning Language planning aims to decompose a goal into sequences of steps (Kaplan and Baldauf, 1997), which is widely used in robotics (Kaiser et al., 2014; Paxton et al., 2019; Berg et al., 2022) and procedural text generation (Goldfarb-Tarrant et al., 2020; Hu et al., 2022). Early studies approach language planning with syntactic parsing for the context (Koller and Stone, 2007; Garoufi and Koller, 2010). Recently, researchers have investigated the planning capability of language models in various domains (Olmo et al., 2021; Valmeekam et al., 2022). However, they mainly focus on generating scripts for stereotypical activities toward abstract goals. For example, Huang et al. (2022) proposes to plan for the general-typed tasks for embodied agents, while Yang et al. (2021) edits actions for abstract goals to video retrieval. In contrast, we explore planning for specific goals (e.g., “*make a cake for diabetics*”). Collins et al. (2022) has benchmarked LLMs for planning with included/excluded objects, but they merely study this problem in a limited scope (only dozens of cases) without further in-depth analysis.

Scripts A structure describing a sequence of events in a particular scenario is *script* (Schank and Abelson, 1975), consisting of two types: 1) *Narrative script*: a narrative chain of events describing a particular scenario derived from narrative texts such as recipes (Fang et al., 2022) or stories (Tandon et al., 2020); 2) *Goal-oriented script* (Regneri et al., 2010; Wanzare et al., 2016): an appropriate sequence of steps as instructions to achieve a goal. In this work, the steps for achieving a given goal in language planning can be categorized into the second class. Many datasets for goal-oriented scripts have been proposed to improve the language planning ability of LMs (Sakaguchi et al., 2021; Lyu et al., 2021). However, they mainly consist of abstract goals with prototypical instructions and thus are not built to train LMs for planning with more specific goals.

Constraint Type 1: Modifier
Definition: A word, an adjective or a phrase that modifies or constrains an abstract goal.
Ex.1: Make a chocolate cake.
Ex.2: Make a pink cake.
Constraint Type 2: Method
Definition: A tool or specified mode that controls the process for achieving the goal.
Ex.1: Make a cake with an oven .
Ex.2: Make a cake by using cake mix .
Constraint Type 3: Intent
Definition: An additional purpose or demand when completing the goal.
Ex.1: Make a cake for wedding .
Ex.2: Make a cake for diabetics .

Table 1: Three types of constraints and their definitions that are used to prompt for new instances of specific goals. In the examples (Ex.), upon the abstract goal, we give two instances for each type of constraint by combining the goal with constraints into specific goals. The constraint within each example is **highlighted**.

In-Context Learning With the great success of LLMs (Brown et al., 2020; Ouyang et al., 2022; Chowdhery et al., 2022), *in-context learning* (Brown et al., 2020; Min et al., 2022) has established its great task-solving potentials with a textual task instruction and a few examples. Moreover, when being used for dataset construction, the data samples that LLMs generate can sometimes outperform crowd-sourced human-authored data in factuality and fluency (Lu et al., 2022a; Min et al., 2022). This shows a promising alternative to costly large-scale crowd-sourcing to construct datasets using LLMs (Wiegrefe et al., 2022; Liu et al., 2022a; West et al., 2022). Inspired by these studies, in our work, we adopt the in-context learning for LLMs not only for better language planning, but also as a reliable *crowd-worker* to scale up the planning data into a reusable dataset to train smaller models.

3 Definitions

Before diving into technical details, we first clarify some important terms used in the paper.

Scripts A goal-oriented *script* is a list of steps ($\mathbf{S} = \{s_1, s_2, \dots, s_{|\mathbf{S}|}\}$) that fulfill a certain goal (\mathcal{G}) (e.g., “make a cake”) (Suddendorf and Corballis, 2007; Schank and Abelson, 2013). The language planning task is defined as $\mathcal{M} : \mathcal{G} \rightarrow \mathbf{S}$, where \mathcal{M} is the planning model.

Goals Different from previous studies that focus mostly on abstract goals with prototypical scripts,

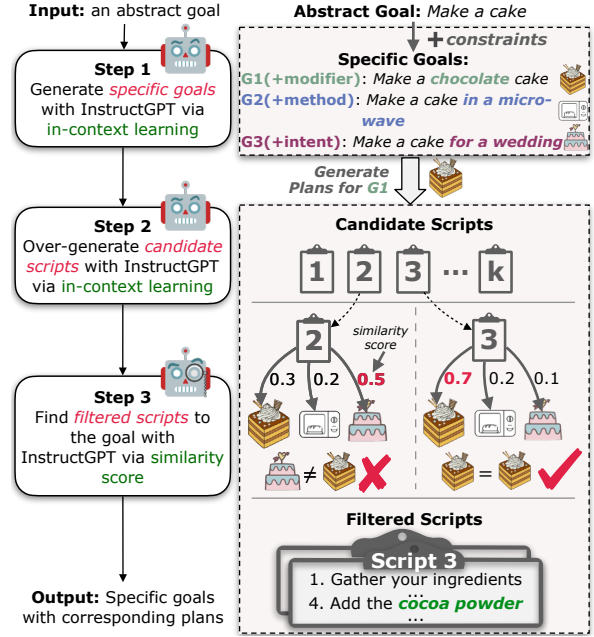


Figure 2: The workflow of using InstructGPT to generate specific goals (Step 1) and planning for the goals with the over-generate-then-filter framework (Step 2-3).

we define a taxonomic structure of goals by extending the derivatives of abstract goals. We define a *specific goal* that inherits from an *abstract one* but with new information as a constraint to limit the scope. An *abstract goal*, denoted as \mathcal{G}_a , refers to stereotypical activities, e.g., “make a cake”. A *specific goal*, denoted as \mathcal{G}_c , is derived from the corresponding \mathcal{G}_a with various constraints, e.g., “make a *chocolate cake*”.

Constraints and Constrained Language Planning To enrich the semantics of specific goals, we define three types of *constraints*, i.e., *modifier*, *method* and *intent*, as shown in Table 1. They express different angles of extending an abstract goal and can be further instantiated and concretized. *Constrained language planning* denotes generating a constraint-faithful script $\mathbf{S} : \mathbf{S} = \mathcal{M}(\mathcal{G}_c)$ toward specific goals (\mathcal{G}_c) with various constraints (\mathcal{C}).

4 Constrained Language Planning with LLMs

In this section, we evaluate and enhance the constraint language planning ability of LLMs. The overall workflow is illustrated in Figure 2. We first extend the specific goals \mathcal{G}_c from the abstract ones \mathcal{G}_a using a human-in-the-loop acquisition approach with LLMs (§ 4.2, Step 1), and propose an over-generate-then-filter framework to

I: Specific Goal Generation
<i>/* Task prompt */</i> Create possible Specific Goals according to the Abstract Goal when the Constraint Type is <i>Modifier</i> . <i>/* Examples */</i> Abstract Goal: Say Goodbye in Different Language Constraint: French; Specific Goal: Say Goodbye in French Constraint: English; Specific Goal: Say Goodbye in English <i>/* Auto completion of constraints and specific goals */</i> Abstract Goal: Make a cake Constraint: <i>Chocolate</i> ; Specific Goal: <i>Make a chocolate cake</i> Constraint: <i>Vanilla</i> ; Specific Goal: <i>Make a vanilla cake</i>
II: Script Generation
<i>/* Task prompt */</i> List the steps of making a cake based on Constraint and Specific Goal. <i>/* Examples */</i> Goal: Make a cake Steps: 1. Gather your ingredients. 2. ... Goal: Make a cupcake Steps: 1. Decide on a pattern. 2. ... <i>/* Auto-completion of script for a specific goal */</i> Constraint: Chocolate; Specific Goal: Make a chocolate cake Steps: <i>1. Gather ingredients. ... 4. Add the cocoa powder...</i>

Table 2: Examples of prompt for InstructGPT for specific goal generation and script generation via in-context learning. Generated texts are *highlighted*.

obtain scripts (§ 4.3, Step 2-3). Then, we reveal that LLMs (e.g., GPT-3 (Brown et al., 2020), InstructGPT (Ouyang et al., 2022)) are prone to be unfaithful to the constraints in \mathcal{G}_c , and our approach can alleviate this problem (§ 4.4). We use `text-davinci-002` as the default InstructGPT variant, which has $\geq 175\text{B}$ parameters.²

4.1 In-Context Learning for LLMs

We deploy LLMs for constrained language planning via in-context learning (Brown et al., 2020; Ouyang et al., 2022). Given a task input (X), we first write a task prompt (T) describing the task, and then provide several examples ($E = \{E_i\}_{i=1}^{|E|}$, where $E_i = (X_i, Y_i)$ are used for few-shot learning). An LLM generates output (Y) by completing the prompt ($Y = \mathcal{M}(T, E, X)$). The whole process does not require any gradient update, allowing LLMs to generate new specific goals and scripts without massive training data.

²Code names and approximated parameters of GPT-3 models are based on <https://blog.eleuther.ai/gpt3-model-sizes/> and <https://beta.openai.com/docs/models>. Note that OpenAI does not release detailed information about later versions of GPT-3, and thus for brevity, we default its size to 175B.

Data Source for Examples We adopt wikiHow (Koupae and Wang, 2018), a data source of instructional articles on various topics, as the initial dataset for providing examples. The articles are titled as “*how to ...?*”, describing abstract goals, and consist of steps to achieve them. We use the titles (\mathcal{G}_a) and steps (\mathcal{S}) as examples.

4.2 Acquisition of Specific Goals

Since no dataset of specific goals exists to support our study, we have to acquire these goals first. As elaborated in Table 1, we extend the abstract goals with multi-faceted constraints for human-in-the-loop data acquisition using InstructGPT.

First, we manually prepare a pool of examples that derive specific goals from an abstract one with constraints.³ Each example is attached to a constraint type (i.e., modifier, method or intent), and contains more than one constraint and specific goal so that InstructGPT is prompted to generate multiple \mathcal{G}_c for one \mathcal{G}_a . Next, given an abstract goal from wikiHow, we enumerate each constraint type to ensure data diversity. Then, we sample several examples of the constraint type from the pool. Finally, we input the task prompt, examples and the \mathcal{G}_a into InstructGPT for the completion of \mathcal{G}_c .

An example in Table 2 (I) shows InstructGPT generates constraints “*chocolate*” and “*vanilla*” for \mathcal{G}_a (“*make a cake*”) given the constraint type *modifier* and some examples, and completes the specific goals (“*make a chocolate cake*” and “*make a vanilla cake*”).

4.3 Acquisition of Scripts

After getting specific goals with constraints, we can test the ability of LLMs to fulfill them.

Planning with InstructGPT We first write a task prompt T . Given the \mathcal{G}_c , we back-trace its \mathcal{G}_a and extract the verbs (“*make*”) and nouns (“*cake*”) from \mathcal{G}_a . Then we use the verbs and nouns as keywords to retrieve two similar goals as examples E from the wikiHow dataset. Finally, the task prompt T , examples E and \mathcal{G}_c with constraint \mathcal{C} are fed into InstructGPT. As shown in Table 2 (II), we adopt the scripts, i.e., “*make a cake*” and “*make a cupcake*” to prompt InstructGPT to generate a script for “*make a chocolate cake*”.

Over-Generation and Filtering Using the above-mentioned approach, generated scripts by

³Complete examples can be found in Appendix B.1.

InstructGPT are reasonable and fluent. However, they sometimes are not faithful to the constraints under closer examination (§ 4.4). Previous studies have shown that the output quality of LLMs falls in high variance (Wiegrefe et al., 2022), leading to bad performance. Thus, we adopt the idea of over-generate-then-filter to improve generation quality, which is shown to be effective in previous work (Wiegrefe et al., 2022; Liu et al., 2022a). We over-generate K sampled from InstructGPT.⁴

Next, a filter model is developed to select the faithful scripts. Due to the diverse expressions of language, we rely not on rules and patterns (*i.e.*, constraint words must appear in the script), but on the semantic similarity between goals and scripts for filtering. For example, “*decorating the cake with candles*” could be a faithful step to make a cake “*for a birthday party*”. Motivated by this, we first collect a set of goals, consisting of the target goal (\mathcal{G}_c^+) as a positive sample and others ($\{\mathcal{G}_c^-\}$) generated from the same abstract goal (\mathcal{G}_a) as negative samples. In the previous case, the negatives include “*make a cake in the microwave*” and “*make a cake for a wedding*”. We convert scripts and goals into InstructGPT embeddings (text-embedding-ada-002) and calculate cosine similarity as similarity scores to measure semantic similarity. Additionally, we reward the script that explicitly contains the keywords of the target constraint. We only keep the script if \mathcal{G}_c^+ scores the highest in the goal set.

4.4 Evaluation

We randomly collect 100 abstract goals (*e.g.*, “*make a cake*”) from `wikiHow` and conduct manual evaluations on the generated specific goals and their scripts. We compare our methods with instruction tuning methods, T0 (Sanh et al., 2022) and Flan-T5 (Chung et al., 2022), vanilla GPT-3 (Ouyang et al., 2022) with different sizes, Codex (Chen et al., 2021) and InstructGPT (Ouyang et al., 2022) with different sizes. We also add “*Let’s think step by step*” before each answer for script generation, which is a simple but effective trick to improve zero-shot reasoning for LLMs (Kojima et al., 2022). For a retrieval baseline, we directly use the goals to search and retrieve the most relevant scripts from

⁴In practice, $K = 2$ is sufficient, as shown in Appendix B.3. Intuitively, the reason this approach works is that the generation accuracy can be improved from $1 - p$ to $1 - p^K$ (at least one is correct), where p is the probability that InstructGPT generates a wrong script.

Model	Modifier	Method	Intent	All
Retrieval	26.67	38.89	35.71	34.00
T0 (11B)	30.00	21.12	25.00	24.00
Flan-T5 (11B)	50.00	42.25	31.25	42.00
GPT-3 (1.3B)	13.33	12.96	18.75	14.00
GPT-3 (6.7B)	23.33	7.40	25.00	15.00
GPT-3 (175B)	30.00	22.22	25.00	25.00
Codex (175B)	46.67	55.56	18.75	47.00
InstructGPT (1.3B)	20.00	22.22	28.57	22.00
InstructGPT (6.7B)	60.00	42.25	43.75	47.00
InstructGPT (175B)	73.33	74.08	42.86	69.00
+ “ <i>let’s think step...</i> ”	70.00	75.92	50.00	68.00
+ Our Method	96.67	98.15	92.86	95.00
w/ $f_{sim} = \text{SBERT}$	86.66	74.89	81.25	78.00
w/ $f_{sim} = \text{SimCSE}$	73.33	78.73	75.00	75.00
w/ $f_{sim} = \text{None}$	93.33	94.44	87.50	93.00

Table 3: Accuracy (%) of generated scripts for different constraint types by manual evaluation. f_{sim} denotes the choice for similarity function during filtering, *i.e.*, replacing InstructGPT embedding with that of SimCSE (Gao et al., 2021) and Sentence-BERT (Reimers and Gurevych, 2019). $f_{sim} = \text{None}$ denotes we only reserve the scripts that contain constraint words.

the `wikiHow` website⁵ as results.

Are specific goals generated by LLMs of high quality? We ask InstructGPT to generate 300 ($3\times$) specific goals for 3 constraint types based on the 100 abstract goals from `wikiHow`. For evaluation, we recruit annotators on Amazon Mechanical Turk to check whether these goals are correct. Each case is examined by three annotators, who reach an inter-rater agreement at Fleiss’s $\kappa = 0.86$ (Fleiss et al., 1981). InstructGPT achieves 98.00% accuracy, indicating that LLMs can derive specific goals of rather high quality.

Can LLMs write scripts for specific goals? To answer this question, we first let InstructGPT generate scripts for the 100 abstract goals from `wikiHow` and ask three annotators to check the correctness of the scripts (with Fleiss’s $\kappa = 0.79$). The correctness is decided by both the fulfillment of the goal and the completeness of the semantics. InstructGPT achieves 97.00% accuracy, proving that LLMs can plan for abstract goals very well. However, it is not the case for specific goals. We sample 100 specific goals from 300 generated ones (mentioned above) and evaluate the scripts generated from baselines and our method.

Table 3 reports the overall accuracy of the results.

⁵<https://www.wikihow.com/Main-Page>

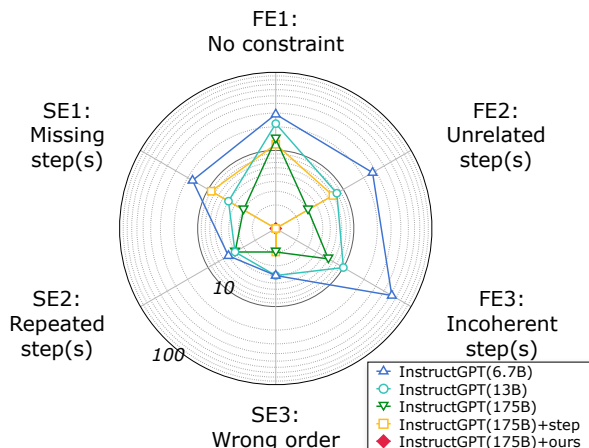


Figure 3: Errors of the generated scripts by human evaluation. The axis of the radar chart is in \log -scale. Notably, ours reduces to virtually one dot in the graphic because it does not have many errors (0-1%). SE and FE denote semantic completeness and faithfulness error.

We find that: 1) Overall, all baselines achieve unsatisfactory results on planning for specific goals, with InstructGPT outperforming others. Especially, the scripts with *intent*-type constraints have the worst accuracy, and adding “*let’s think step-by-step*” does not help much; 2) The retrieval from *wikiHow* does not lead to the desired script; 3) With our method, InstructGPT can generate scripts of higher quality by a large margin; 4) Replacing the similarity function with embeddings from other pre-trained models results in performance drops.

What types of errors do LLMs usually make in this task? To respond to the motivations of our methods, we conduct detailed analyses to investigate why LLMs fail. We evaluate the model planning performance in two aspects: 1) *Semantic completeness* (SE): whether the steps in the script are missing, repeated or in the wrong order; 2) *Faithfulness to the constraints* (FE): whether the script is faithful to the constraints and the steps are coherent (related) within the script. We define six types of errors upon the two, *i.e.*, *i*) SE: missing, repeated step(s) and wrong order and *ii*) FE: no constraint, unrelated step(s) or incoherent step(s).⁶ Annotators are asked to review 100 scripts generated by InstructGPT and mark the error types.⁷ Results in Figure 3 show that: 1) The semantic completeness in generated scripts is acceptable, but the faithfulness to the constraints can not be guaran-

⁶The detailed definitions can be found in Appendix B.4.

⁷The case study of how InstructGPT fails at planning for specific goals is shown in Appendix B.5.

Work	0.0	0.67	1.0	1.0	1.0
Sports	0.25	0.5	1.0	1.0	1.0
Relationships	0.0	0.5	1.0	1.0	1.0
Personal Care	0.14	0.33	0.5	0.6	1.0
Home	0.17	0.3	0.5	0.6	1.0
Traditions	0.25	0.4	0.5	0.57	1.0
Hobbies	0.1	0.42	0.6	0.67	0.88
Health	0.33	0.5	0.57	0.75	0.9
Food	0.28	0.66	0.78	0.78	0.94
Business	0.33	0.5	1.0	0.67	1.0
Family Life	0.0	0.0	1.0	0.67	1.0
Education	0.2	0.33	0.5	0.67	1.0
Electronics	0.25	0.38	0.6	0.62	0.92
Vehicles	0.0	0.0	0.75	0.5	1.0
Arts	0.33	0.29	1.0	0.5	1.0
	6.7B	13B	175B	Step	Ours

Figure 4: The heat-map depicts the human-evaluated script accuracy of different methods in different topic categories for specific goals.

teed; 2) Our method greatly improves the planning quality both in semantic completeness and faithfulness.

What kinds of goals do InstructGPT typically fail? By far, we already know that LLMs fail at specific goals, especially for *intent*-type constraints. We dig into more fine-grained topic categories of constraints defined in *wikiHow*. The heat map in Figure 4 shows that the planning performance of InstructGPTs varies considerably for goals of different categories, and the planning accuracy for each category improves greatly with our method.

5 Script Distillation from LLMs

Since LLMs are costly to deploy, it is essential to enable language planning ability for smaller, specialized models. Creating datasets is an inevitable step to this end. However, previous datasets do not enable planning for specific goals (Sakaguchi et al., 2021; Lyu et al., 2021), and manual dataset annotation is expensive and highly demanding. Thus, we follow the idea of *symbolic knowledge distillation* (West et al., 2022) to distill constrained language planning datasets from LLMs.

5.1 CoScript: A Dataset for Constrained Language Planning

We now apply our method for building a first-of-its-kind **Constrained Script** dataset of language planning, named as **CoScript**. Experiments in § 4 show that LLMs can generate high-quality specific goals and scripts with our over-generating-then-filter framework. We now scale up the experiments

Dataset	# Size	# UT	Avg _{G_c} #	Avg _S #
proScript	6,414	8,826	0	5.45
wikiHow	112,111	158,117	0.42	5.93
CoScript	55,000	76,317	4.13	5.96

Table 4: Statistics of CoScript and previous script datasets proScript and wikiHow, w.r.t. data size, number of unique tokens (# UT), the average number of specific goals for each abstract ones (Avg_{G_c} #), and the average number of steps in scripts (Avg_S #).

for a large-scale dataset. We collect 14,945 article titles as seed abstract goals and retrieve 34,260 similar goals with scripts from wikiHow as examples to prompt InstructGPT (175B) for data generation. Following § 4, dataset construction process consists of three steps, as in Figure 2: 1) We first enumerate constraint types with examples for InstructGPT and obtain specific goals (after de-duplication) based on the seed abstract goals. 2) Then, InstructGPT over-generates K scripts for the specific goals and 3) our filter framework selects the faithful scripts as the final data.⁸

In total, we generate 55,000 specific goals with corresponding scripts. We randomly choose 2,000 data as the validation set and 3,000 data as the test set. To ensure the quality of the validation and test set, we ask crowd-sourced workers to find and revise the incorrect samples. By collecting the annotation data for error identification of these 5,000 samples, we estimate to achieve 97.80% accuracy for specific goals and 94.98% for constrained script generation, consistent with the results in Table 3.

5.2 Dataset Analysis

Script Diversity Analysis As shown in Table 4, despite the larger scale of wikiHow, CoScript has more specific goals than wikiHow and thus is valuable for the constrained language planning task. Besides, previous studies (Fu et al., 2021; Narayan et al., 2022) find that the texts generated by LMs may be too repetitive and less diverse. For this concern, we compare our CoScript with a recent goal-oriented script dataset proScript (Sakaguchi et al., 2021) created by crowd-sourcing. As reported in Table 4, 1) CoScript is much larger than proScript, with more scripts and a higher number of steps per script; 2) CoScript exhibits high lexical diversity, with more unique words than human-written proScript.

⁸Details about hyper-parameters and costs can be found in Appendix C.1.




Figure 5: Statistics of constraint types in CoScript dataset, with representative topic categories or the first words for each constraint type.

Constraint Analysis Figure 5 shows the constraint distribution of CoScript. We compute the proportions of constraint types with their representative categories obtained from Probase (Wu et al., 2012), and the initial words of constraint instances. We find CoScript shows high heterogeneity and pluralism in the generated specific goals. Interestingly, InstructGPT tends to start with the word “if” or “when” for hypothetical constraints (e.g., “if someone is lactose intolerant” for “make a cake”), suggesting the potential for future research on counterfactual reasoning in language planning. We also analyze the domain distribution of CoScript in the Appendix C.2

6 Constrained Language Planning with Specialized Models

With CoScript, we can train smaller but specialized models for constrained language planning.

6.1 Experimental setup

Baselines We use GPT-2 (causal LM) (Radford et al., 2019) and T5 (encoder-decoder LM) (Raffel et al., 2020) as baselines. Given goals, the models are trained to generate a list of steps S for planning. Moreover, we adopt the idea of retrieval-augmented text generation (Lewis et al., 2020) and add retrieved examples in the input to improve generation quality.

Metrics We use BLEU (Papineni et al., 2002), ROUGE-L (Lin, 2004) and BERTScore (Zhang et al., 2020) as automatic metrics to measure semantic completeness. We also train a binary classification model to decide whether the generated texts are faithful to the constraints. Specifically, we

Model	Faithful	ROUGE	BLEU	BERTScore
<i>Trained on wikiHow</i>				
GPT-2	64.93	20.28	17.91	80.74
GPT-2 (large)	62.20	23.74	24.69	83.63

T5 (base)	86.13	20.30	15.48	79.02
T5 (large)	85.13	22.95	20.60	82.27
T5 (3B)	77.90	20.72	16.95	81.01

<i>Trained on CoScript</i>				
GPT-2	74.60	28.09	26.75	84.72
GPT-2 (large)	76.73	30.60	30.22	85.77
+retrieval	76.30	32.78	32.92	86.41

T5 (base)	91.53	26.53	22.06	83.14
T5 (large)	91.87	29.40	29.14	83.48
+retrieval	86.03	35.91	36.10	87.39
T5 (3B)	93.00	45.68	43.83	90.18
+retrieval	92.53	46.54	47.62	90.84

Table 5: Overall script generation performance for models trained on different training sets. Note that the test set is the same for all models.

Model	Modifier	Method	Intent	All
T5 (large)	91.54	92.57	90.21	91.81
+retrieval	87.39	85.86	84.44	86.03

GPT-2 (large)	78.78	78.77	69.48	76.73
+retrieval	77.33	78.28	70.97	76.30

Table 6: **Faithfulness** scores of specialized models for each constraint type on the test set of CoScript.

collect 50,000 data from CoScript as positive samples, and shuffle the goals and scripts to construct 50,000 negative ones. Then, we fine-tune a DeBERTa (v3 large) model (Khashabi et al., 2020) for classification, achieving 91.53% accuracy on the test set.

Training Data To gain a fine-grained perspective on planning toward specific goals, we train LMs on both wikiHow (\mathbb{D}_{tr}^{wi}) and CoScript (\mathbb{D}_{tr}^{co}), and test them on CoScript test set (\mathbb{D}_{te}^{co}). Both datasets share *similar* scripts, but the goals in wikiHow are mostly abstract ones. For wikiHow, we also randomly collect 50,000 goals with scripts as \mathbb{D}_{tr}^{wi} .

6.2 Results

The comparison for models trained on wikiHow and CoScript are shown in Table 5. In general, LMs trained on CoScript outperform that on wikiHow. T5 outperforms GPT-2 in faithfulness, possibly due to its encoder-decoder framework being better at handling input information. However, GPT-2 outperforms T5 on other text generation metrics for scripts. This could be because CoScript




Figure 6: The faithfulness curves when altering the proportions of CoScript (α) and wikiHow ($1 - \alpha$) in a fixed-size training set.

Model	Size	Modifier	Method	Intent	All
GPT-3	175B	30.00	22.22	25.00	25.00
Codex	175B	46.67	55.56	18.75	47.00
InstructGPT	175B	73.33	74.08	42.86	69.00

T5 (wikiHow)	3B	20.00	12.96	6.25	14.00
T5 (CoScript)	3B	63.33	55.55	43.75	56.00
+retrieval	3B	76.66	66.66	75.00	71.00

Table 7: **Accuracy** (%) of scripts generated by different models. We fine-tune a T5 (3B) on wikiHow and CoScript while deploying LLMs via few-shot in-context learning.

is distilled from InstructGPT, leading to a biased data distribution that favors decoder-only causal language models, *e.g.*, the GPT family.

Based on Table 5, we find that augmenting models with retrieved examples can improve semantic completeness. However, the constraint faithfulness could be undermined as models tend to mimic the retrieved examples. To further understand the role of retrieval augmentation, we conduct a manual evaluation that based on 100 random samples generated by T5 (3B) with and without retrieval augmentation. We discover that 57% of T5’s results are correct, and the number goes up to 70% with retrieval augmentation. Thus, although we observe a slight drop in faithfulness score (93.00 \rightarrow 92.53 from Table 5), retrieval augmentation still brings much improvement over the base model.

Faithfulness of Constraints of Different Types Will LLMs’ planning preferences for constraint types pass to the specialized models? We find the results in Table 6 are consistent with that of LLMs (Table 3). Specialized models are also the worst at specific goals with intent-typed constraints.

CoScript vs. wikiHow We mix two datasets together with a hyper-parameter α to control the proportion of two datasets, where the new training

set $\mathbb{D}_{\text{tr}} = \alpha \mathbb{D}_{\text{tr}}^{\text{co}} + (1 - \alpha) \mathbb{D}_{\text{tr}}^{\text{wi}}$. By altering α (constant data size), the faithfulness curves in Figure 6 shows that adding more data from CoScript consistently improves model performance in constraint faithfulness. Thus, training on CoScript contributes to more faithful planners.

Specialized Models vs. LLMs We further fine-tune a T5 (3B) on CoScript and wikiHow to generate scripts for the specific goals in § 4.4, which are held out from the training set. Table 7 shows that T5 fine-tuned on CoScript with retrieval augmentation can generate scripts of higher quality than most LLMs in Table 3, indicating that smaller models can surpass larger models when properly trained on suitable datasets.

7 Conclusion

In this paper, we define planning toward specific goals with constraints. We propose a better prompting method for LLMs, and distill a novel dataset from LLMs (CoScript) to improve the constrained language planning ability of specialized models. Experiments show that our method improves the planning quality of LLMs for specific goals, and smaller models trained on CoScript even outperform LLMs. We hope the CoScript dataset will be a valuable resource to advance the research on language planning with more complex and diverse goals and constraints.

Limitations

The proposed method for improving LLMs is a post-hoc re-ranking approach, and we do not improve LLMs themselves due to the difficulty of fine-tuning LLMs. Besides, we improve the ability of constrained language planning for smaller models from the perspective of building task-related datasets, but do not consider investigating the model itself, other than adopting retrieval augmentation. In addition, because automatic metrics for generated text are limited, the automatic evaluation of this paper may result in an overestimation or underestimation of the mentioned methods, though we attempt to mitigate this by incorporating a moderate amount of human evaluation. Despite the advanced planning capabilities of newer language models, our work remains significantly valuable to the knowledge distillation of these LLMs into smaller and more cost-effective models.

We also discover several limitations of the proposed CoScript datasets. First, the specific goal

explored in this work only inherits from an abstract one with one extra constraint. However, in real-life situations, complex planning may involve multiple constraints, which we do not investigate in this work. Another limitation of CoScript is that our dataset is generated from InstructGPT, and thus the data distributions may be biased to favor causal language models. This is a common issue with machine-generated datasets, which we address by manually curating CoScript’s validation and test sets. Furthermore, there are still some incorrect samples (about 5%) in the training data without manual correction due to the limits of budget and time. Last but not least, we only consider whether the script can be executed at the human level. The script execution for robots (Huang et al., 2022; Lu et al., 2022b) is unstudied in our work, and there still exist huge gaps in transferring complex human language to one that is understandable and executable by robots.

Ethics Statement

Use of Human Annotations We protect the privacy rights of crowd-sourced workers and pay them above the local minimum wage. We use Amazon Mechanical Turk (AMT) and require 300 annotators to be located in the U.S. as a proxy for English competency. We pay at a rate of \$6/hour for 20 samples. We acknowledge that constructing datasets from large language models may suffer from toxic language and cause severe risks for social society (Ousidhoum et al., 2021; Baldini et al., 2022). Therefore, we ask the annotators to discard the offensive and harmful data when reviewing the CoScript. However, there may still be prejudicial data in our final dataset that goes unnoticed.

wikiHow Source The content available on wikiHow is shared under a Creative Commons License (CC-BY-NC-SA)⁹, which permits others to share, copy, distribute, and adapt the content for non-commercial purposes. In our research, we use wikiHow as an initial dataset for providing examples to construct our dataset. Our dataset is released on GitHub and is only used to advance academic research on language planning with more complex and diverse goals and constraints. Therefore, we emphasize that our usage aligns with the requirements under the license.

⁹<https://creativecommons.org/licenses/by-nc-sa/3.0/>

Covered Domains in CoScript CoScript is derived from wikiHow and encompasses 19 daily life goal categories (as illustrated in Figure 8). These categories cover a wide range of practical topics of everyday life. However, as shown in Figure 8, we emphasize that sensitive and high-risk domains, including medical, legal, and high-stakes financial advice, are excluded from the dataset to minimize potential risks related to inaccurate or misleading information. We encourage researchers and developers to leverage this dataset to build models that accurately understand and respond to user queries on various non-sensitive, non-critical topics.

Factuality, Toxicity and Biases We recognize that the factuality of generated content is crucial, especially in high-stakes scenarios. Therefore, annotators are asked to verify the consistency between generated scripts and goals with constraints for validation and test sets. They also assess and revise the content to minimize hallucinations, factual errors, and any inappropriate or misleading information.

Previous work found that LLMs may generate toxic contents (Cao et al., 2022; Liu et al., 2022b). We highlight that our dataset is not intended for safety-critical applications or as a substitute for expert advice in such domains. Annotators are specifically instructed to discard offensive and harmful data during the review of the validation and test sets in CoScript. However, despite these precautions, there may still be some prejudicial data that goes unnoticed in our final dataset.

Acknowledgement

We thank the anonymous reviewers for their valuable comments, and Wei Shi and Shuang Li from Fudan University for their useful suggestions for the manuscript. This work is supported by the Chinese NSF Major Research Plan (No.92270121), Shanghai Science and Technology Innovation Action Plan (No.21511100401) and the Science and Technology Commission of Shanghai Municipality Grant (No. 22511105902).

References

Robert P Abelson. 1976. Script processing in attitude formation and decision making.

Ioana Baldini, Dennis Wei, Karthikeyan Natesan Ramamurthy, Moninder Singh, and Mikhail Yurochkin. 2022. [Your fairness may vary: Pretrained language](#)

[model fairness in toxic text classification](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 2245–2262, Dublin, Ireland. Association for Computational Linguistics.

Matthew Berg, George Konidaris, and Stefanie Tellex. 2022. [Using language to generate state abstractions for long-range planning in outdoor environments](#). In *2022 International Conference on Robotics and Automation (ICRA)*, pages 1888–1895.

Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. 2020. [Language models are few-shot learners](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 1877–1901. Curran Associates, Inc.

Meng Cao, Yue Dong, and Jackie Cheung. 2022. [Hallucinated but factual! inspecting the factuality of hallucinations in abstractive summarization](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 3340–3354, Dublin, Ireland. Association for Computational Linguistics.

Mark Chen, Jerry Tworek, Heewoo Jun, Qiming Yuan, Henrique Ponde de Oliveira Pinto, Jared Kaplan, Harri Edwards, Yuri Burda, Nicholas Joseph, Greg Brockman, et al. 2021. [Evaluating large language models trained on code](#). *arXiv preprint arXiv:2107.03374*.

Aakanksha Chowdhery, Sharan Narang, Jacob Devlin, Maarten Bosma, Gaurav Mishra, Adam Roberts, Paul Barham, Hyung Won Chung, Charles Sutton, Sebastian Gehrmann, et al. 2022. [Palm: Scaling language modeling with pathways](#). *arXiv preprint arXiv:2204.02311*.

Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Eric Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, et al. 2022. [Scaling instruction-finetuned language models](#). *arXiv preprint arXiv:2210.11416*.

Katherine M Collins, Catherine Wong, Jiahai Feng, Megan Wei, and Joshua B Tenenbaum. 2022. [Structured, flexible, and robust: benchmarking and improving large language models towards more human-like behavior in out-of-distribution reasoning tasks](#). *arXiv preprint arXiv:2205.05718*.

Biaoyan Fang, Timothy Baldwin, and Karin Verspoor. 2022. [What does it take to bake a cake? the RecipeRef corpus and anaphora resolution in procedural text](#). In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 3481–3495,

- Dublin, Ireland. Association for Computational Linguistics.
- Joseph L Fleiss, Bruce Levin, Myunghee Cho Paik, et al. 1981. The measurement of interrater agreement. *Statistical methods for rates and proportions*, 2(212-236):22–23.
- Zihao Fu, Wai Lam, Anthony Man-Cho So, and Bei Shi. 2021. [A theoretical analysis of the repetition problem in text generation](#). *Proceedings of the AAAI Conference on Artificial Intelligence*, 35(14):12848–12856.
- Tianyu Gao, Xingcheng Yao, and Danqi Chen. 2021. [SimCSE: Simple contrastive learning of sentence embeddings](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 6894–6910, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Konstantina Garoufi and Alexander Koller. 2010. [Automated planning for situated natural language generation](#). In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, pages 1573–1582, Uppsala, Sweden. Association for Computational Linguistics.
- Seraphina Goldfarb-Tarrant, Tuhin Chakrabarty, Ralph Weischedel, and Nanyun Peng. 2020. [Content planning for neural story generation with aristotelian rescoring](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4319–4338, Online. Association for Computational Linguistics.
- Zhe Hu, Hou Pong Chan, Jiachen Liu, Xinyan Xiao, Hua Wu, and Lifu Huang. 2022. [PLANET: Dynamic content planning in autoregressive transformers for long-form text generation](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 2288–2305, Dublin, Ireland. Association for Computational Linguistics.
- Wenlong Huang, Pieter Abbeel, Deepak Pathak, and Igor Mordatch. 2022. [Language models as zero-shot planners: Extracting actionable knowledge for embodied agents](#). In *International Conference on Machine Learning, ICML 2022, 17-23 July 2022, Baltimore, Maryland, USA*, volume 162 of *Proceedings of Machine Learning Research*, pages 9118–9147. PMLR.
- Peter Kaiser, Mike Lewis, Ronald P. A. Petrick, Tamim Asfour, and Mark Steedman. 2014. [Extracting common sense knowledge from text for robot planning](#). In *2014 IEEE International Conference on Robotics and Automation (ICRA)*, pages 3749–3756.
- Robert B Kaplan and Richard B Baldauf. 1997. *Language planning from practice to theory*, volume 108. Multilingual Matters.
- Daniel Khashabi, Sewon Min, Tushar Khot, Ashish Sabharwal, Oyvind Tafjord, Peter Clark, and Hananeh Hajishirzi. 2020. [UNIFIEDQA: Crossing format boundaries with a single QA system](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1896–1907, Online. Association for Computational Linguistics.
- Takeshi Kojima, Shixiang Shane Gu, Machel Reid, Yutaka Matsuo, and Yusuke Iwasawa. 2022. [Large language models are zero-shot reasoners](#). In *ICML 2022 Workshop on Knowledge Retrieval and Language Models*.
- Alexander Koller and Matthew Stone. 2007. [Sentence generation as a planning problem](#). In *Proceedings of the 45th Annual Meeting of the Association of Computational Linguistics*, pages 336–343, Prague, Czech Republic. Association for Computational Linguistics.
- Mahnaz Koupaee and William Yang Wang. 2018. [Wikihow: A large scale text summarization dataset](#). *ArXiv preprint*, abs/1810.09305.
- Alexander Kovalchuk, Shashank Shekhar, and Ronen I. Brafman. 2021. [Verifying plans and scripts for robotics tasks using performance level profiles](#). *Proceedings of the International Conference on Automated Planning and Scheduling*, 31(1):673–681.
- Patrick Lewis, Ethan Perez, Aleksandra Piktus, Fabio Petroni, Vladimir Karpukhin, Naman Goyal, Heinrich Küttler, Mike Lewis, Wen-tau Yih, Tim Rocktäschel, Sebastian Riedel, and Douwe Kiela. 2020. [Retrieval-augmented generation for knowledge-intensive nlp tasks](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 9459–9474. Curran Associates, Inc.
- Chin-Yew Lin. 2004. [ROUGE: A package for automatic evaluation of summaries](#). In *Text Summarization Branches Out*, pages 74–81, Barcelona, Spain. Association for Computational Linguistics.
- Alisa Liu, Swabha Swayamdipta, Noah A. Smith, and Yejin Choi. 2022a. [WANLI: Worker and AI collaboration for natural language inference dataset creation](#). In *Findings of the Association for Computational Linguistics: EMNLP 2022*, pages 6826–6847, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Tianyu Liu, Yizhe Zhang, Chris Brockett, Yi Mao, Zhifang Sui, Weizhu Chen, and Bill Dolan. 2022b. [A token-level reference-free hallucination detection benchmark for free-form text generation](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6723–6737, Dublin, Ireland. Association for Computational Linguistics.
- Yao Lu, Max Bartolo, Alastair Moore, Sebastian Riedel, and Pontus Stenetorp. 2022a. [Fantastically ordered prompts and where to find them: Overcoming few-shot prompt order sensitivity](#). In *Proceedings of the*

- 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers), pages 8086–8098, Dublin, Ireland. Association for Computational Linguistics.
- Yujie Lu, Weixi Feng, Wanrong Zhu, Wenda Xu, Xin Eric Wang, Miguel Eckstein, and William Yang Wang. 2022b. Neuro-symbolic causal language planning with commonsense prompting. *arXiv preprint arXiv:2206.02928*.
- Qing Lyu, Li Zhang, and Chris Callison-Burch. 2021. **Goal-oriented script construction**. In *Proceedings of the 14th International Conference on Natural Language Generation*, pages 184–200, Aberdeen, Scotland, UK. Association for Computational Linguistics.
- Sewon Min, Xinxu Lyu, Ari Holtzman, Mikel Artetxe, Mike Lewis, Hannaneh Hajishirzi, and Luke Zettlemoyer. 2022. **Rethinking the role of demonstrations: What makes in-context learning work?** *ArXiv preprint*, abs/2202.12837.
- Swaroop Mishra, Daniel Khashabi, Chitta Baral, Yejin Choi, and Hannaneh Hajishirzi. 2022. **Reframing instructional prompts to GPTk’s language**. In *Findings of the Association for Computational Linguistics: ACL 2022*, pages 589–612, Dublin, Ireland. Association for Computational Linguistics.
- Shashi Narayan, Gonçalo Simões, Yao Zhao, Joshua Maynez, Dipanjan Das, Michael Collins, and Mirella Lapata. 2022. **A well-composed text is half done! composition sampling for diverse conditional generation**. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1319–1339, Dublin, Ireland. Association for Computational Linguistics.
- Alberto Olmo, Sarath Sreedharan, and Subbarao Kambhampati. 2021. Gpt3-to-plan: Extracting plans from text using gpt-3. *FinPlan 2021*, page 24.
- Nedjma Ousidhoum, Xinran Zhao, Tianqing Fang, Yangqiu Song, and Dit-Yan Yeung. 2021. **Probing toxic content in large pre-trained language models**. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 4262–4274, Online. Association for Computational Linguistics.
- Long Ouyang, Jeff Wu, Xu Jiang, Diogo Almeida, Carroll L Wainwright, Pamela Mishkin, Chong Zhang, Sandhini Agarwal, Katarina Slama, Alex Ray, et al. 2022. Training language models to follow instructions with human feedback. *arXiv preprint arXiv:2203.02155*.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. **Bleu: a method for automatic evaluation of machine translation**. In *Proceedings of the 40th Annual Meeting of the Association for Computational Linguistics*, pages 311–318, Philadelphia, Pennsylvania, USA. Association for Computational Linguistics.
- Chris Paxton, Yonatan Bisk, Jesse Thomason, Arunkumar Byravan, and Dieter Foxl. 2019. **Prospection: Interpretable plans from language by predicting the future**. In *2019 International Conference on Robotics and Automation (ICRA)*, pages 6942–6948.
- 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.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, Peter J Liu, et al. 2020. Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 21(140):1–67.
- Michaela Regneri, Alexander Koller, and Manfred Pinkal. 2010. **Learning script knowledge with web experiments**. In *Proceedings of the 48th Annual Meeting of the Association for Computational Linguistics*, pages 979–988, Uppsala, Sweden. Association for Computational Linguistics.
- Nils Reimers and Iryna Gurevych. 2019. **Sentence-BERT: Sentence embeddings using Siamese BERT-networks**. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3982–3992, Hong Kong, China. Association for Computational Linguistics.
- Keisuke Sakaguchi, Chandra Bhagavatula, Ronan Le Bras, Niket Tandon, Peter Clark, and Yejin Choi. 2021. **proScript: Partially ordered scripts generation**. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 2138–2149, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Abhilasha Sancheti and Rachel Rudinger. 2022. **What do large language models learn about scripts?** In *Proceedings of the 11th Joint Conference on Lexical and Computational Semantics*, pages 1–11, Seattle, Washington. Association for Computational Linguistics.
- Victor Sanh, Albert Webson, Colin Raffel, Stephen Bach, Lintang Sutawika, Zaid Alyafeai, Antoine Chaffin, Arnaud Stiegler, Arun Raja, Manan Dey, M Saiful Bari, Canwen Xu, Urmish Thakker, Shanya Sharma Sharma, Eliza Szczechla, Taewoon Kim, Gunjan Chhablani, Nihal Nayak, Debajyoti Datta, Jonathan Chang, Mike Tian-Jian Jiang, Han Wang, Matteo Manica, Sheng Shen, Zheng Xin Yong, Harshit Pandey, Rachel Bawden, Thomas Wang, Trishala Neeraj, Jos Rozen, Abheesht Sharma, Andrea Santilli, Thibault Fevry, Jason Alan Fries, Ryan Teehan, Teven Le Scao, Stella Biderman, Leo Gao,

- Thomas Wolf, and Alexander M Rush. 2022. [Multi-task prompted training enables zero-shot task generalization](#). In *International Conference on Learning Representations*.
- Roger C. Schank and Robert P. Abelson. 1975. Scripts, plans, and knowledge. In *Proceedings of the 4th International Joint Conference on Artificial Intelligence - Volume 1, IJCAI'75*, page 151–157, San Francisco, CA, USA. Morgan Kaufmann Publishers Inc.
- Roger C Schank and Robert P Abelson. 2013. *Scripts, plans, goals, and understanding: An inquiry into human knowledge structures*. Psychology Press.
- Thomas Suddendorf and Michael C Corballis. 2007. The evolution of foresight: What is mental time travel, and is it unique to humans? *Behavioral and brain sciences*, 30(3):299–313.
- Niket Tandon, Keisuke Sakaguchi, Bhavana Dalvi, Dheeraj Rajagopal, Peter Clark, Michal Guerquin, Kyle Richardson, and Eduard Hovy. 2020. [A dataset for tracking entities in open domain procedural text](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6408–6417, Online. Association for Computational Linguistics.
- Karthik Valmeekam, Alberto Olmo, Sarath Sreedharan, and Subbarao Kambhampati. 2022. [Large language models still can't plan \(a benchmark for LLMs on planning and reasoning about change\)](#). In *NeurIPS 2022 Foundation Models for Decision Making Workshop*.
- Xuezhi Wang, Jason Wei, Dale Schuurmans, Quoc Le, Ed Chi, and Denny Zhou. 2022. Self-consistency improves chain of thought reasoning in language models. *arXiv preprint arXiv:2203.11171*.
- Lilian D. A. Wanzare, Alessandra Zarcone, Stefan Thater, and Manfred Pinkal. 2016. [A crowdsourced database of event sequence descriptions for the acquisition of high-quality script knowledge](#). In *Proceedings of the Tenth International Conference on Language Resources and Evaluation (LREC'16)*, pages 3494–3501, Portorož, Slovenia. European Language Resources Association (ELRA).
- Jason Wei, Xuezhi Wang, Dale Schuurmans, Maarten Bosma, brian ichter, Fei Xia, Ed H. Chi, Quoc V Le, and Denny Zhou. 2022. [Chain of thought prompting elicits reasoning in large language models](#). In *Advances in Neural Information Processing Systems*.
- Peter West, Chandra Bhagavatula, Jack Hessel, Jena Hwang, Liwei Jiang, Ronan Le Bras, Ximing Lu, Sean Welleck, and Yejin Choi. 2022. [Symbolic knowledge distillation: from general language models to commonsense models](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4602–4625, Seattle, United States. Association for Computational Linguistics.
- Sarah Wiegrefe, Jack Hessel, Swabha Swayamdipta, Mark Riedl, and Yejin Choi. 2022. [Reframing human-AI collaboration for generating free-text explanations](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 632–658, Seattle, United States. Association for Computational Linguistics.
- Wentao Wu, Hongsong Li, Haixun Wang, and Kenny Q. Zhu. 2012. [Probase: A probabilistic taxonomy for text understanding](#). In *Proceedings of the 2012 ACM SIGMOD International Conference on Management of Data, SIGMOD '12*, page 481–492, New York, NY, USA. Association for Computing Machinery.
- Yue Yang, Joongwon Kim, Artemis Panagopoulou, Mark Yatskar, and Chris Callison-Burch. 2021. [Induce, edit, retrieve: Language grounded multimodal schema for instructional video retrieval](#). *ArXiv preprint*, abs/2111.09276.
- Tianyi Zhang, Varsha Kishore, Felix Wu, Kilian Q. Weinberger, and Yoav Artzi. 2020. [Bertscore: Evaluating text generation with BERT](#). In *8th International Conference on Learning Representations, ICLR 2020, Addis Ababa, Ethiopia, April 26-30, 2020*. OpenReview.net.

A Author Contributions

Siyu Yuan Lead the project, develop the original method and original code, lead the curation of the dataset, contribute to the original experiments, and contribute to the original manuscript.

Jiangjie Chen Conceptualization of the original idea, supervision of the research activity planning and execution, contribution to the original manuscript and figures, and acquisition of financial support for the project.

Ziquan Fu Contribute to the original idea, provide financial support for the project, revise the manuscript, and contribute to data curation.

Xuyang Ge Contribute to the experiments and data curation.

Soham Shah Provide financial support for the project and revise the manuscript.

Charles Robert Jankowski Provide financial support for the project and revise the manuscript.

Yanghua Xiao Provide financial support for the project and revise the manuscript.

Deqing Yang Provide financial support for the project, revise the manuscript, and oversee the research activity execution.

B Implementation Details

B.1 Handpicked Examples for Specific Goal Generation

We follow the instructions proposed by Mishra et al. (2022) to better construct the three examples for specific goal generation. As shown in Table 13, we turn long descriptions into bulleted lists in the task prompt for better generation. In addition, in each example, we list two specific goals with constraints, which can prompt InstructGPT to generate multiple specific goals for the given abstract goals. In our experiment, we conduct the specific goals generation with different numbers of examples and report the accuracy and the total number of generated specific goals for 100 abstract goals. The results in Table 8 show that three examples are good in our settings.

# Examples	Accuracy	# Total
2	95.16%	545
3	96.99%	537
4	95.44%	539

Table 8: The specific goals generation performance of InstructGPT with different numbers of examples.

B.2 Prompts Format

To explore the prompt formats on script generation, we test 100 samples mentioned in § 4.4 without using the task prompt or replacing the original words with other words. The results in Table 9 show that: 1) task prompt can help InstructGPT to better understand the task and thus can improve the model performance on script generation; 2) adopting popular words to write the prompts can better improve the effect of prompts.

Format	Goal	Script
Our Method	98.00	69.00
<i>w/o Task Prompt</i>	94.67	64.00
<i>r. Goal→Scenario</i>	96.00	-
<i>r. Abstract→General</i>	96.67	-
<i>r. Step→Event</i>	-	67.00

Table 9: Accuracy (%) of different prompt formats by manual evaluation. We replace (r.) the words in Table 2 with other words for comparison.

B.3 Over-generation Hyper-parameters

To evaluate the LLMs on planning for specific goals, we randomly sample 100 specific goals and




Figure 7: Statistics of the constraints evaluation results. Human annotators are asked to evaluate K generations for each specific goal on 100 samples. The legends indicate the number of constraint-faithful scripts in K over-generation. We report the proportion of generated scripts faithful to the constraints.

then generate scripts from the baseline and our method. Figure 7 reports the faithfulness to constraints. We find that: 1) the output quality of InstructGPT falls in high variance, and the script may be unfaithful to the constraints; 2) over-generation can amplify the likelihood of constraint satisfaction, and $K = 2$ is sufficient.

B.4 Error Types

As shown in Figure 10, we evaluate the model planning performance on two aspects, *i.e.*, *Semantic completeness* and *Faithfulness to the constraints* (§ 4.4), and define six types of errors.

B.5 Case Study

Table 14 lists three examples by InstructGPT (175B) and our approach. The first and second examples show that the scripts generated by InstructGPT may fail in unfaithfulness to the constraints. The third examples demonstrate that although the scripts generated by InstructGPT can be faithful to the constraints, they may suffer from other error types. In contrast, the over-generating-and-filtering method can amplify the likelihood of high-quality generations and thus can make LLMs better planners.

C CoScript Details

C.1 Generation Hyper-parameters

We queried the `text-davinci-002` model through the OpenAI API on June 25 to July 5, 2022. CoScript is generated under a specified license that is compatible with the conditions under OpenAI API. In total, the generation for CoScript costs about \$5,000. The hyper-parameters for script generation are shown in Table 11. If both generations are faithful, we randomly select one into the dataset.

Aspects	Error Types	Explanation	Example: Make a vanilla cake Constraint: Vanilla	
Semantic Completeness	Wrong order	Steps that are in the wrong order.	Correct Script: 1. Gather your ingredients. 2. Preheat the oven to 325° F and grease and flour a cake pan. 3. Cream the butter and suger. 4. Add the eggs and vanilla. 5. Stir in the cake flour. 6. Pour the batter into the pan. 7. Bake the cake for 1 hour 15 minutes.	Generated Script: 1. Preheat the oven to 325° F and grease and flour a cake pan. 2. Gather your ingredients. 3. Buy your ingredients. 4. Cream the butter and salt. 5. Stir in the cake flour. 6. Have a shower. 7. Pour the batter into the pan. 8. Bake the cake for 1 hour 15 minutes.
	Repeat steps	Steps that are repeated in the script.		
	Missing steps	Important steps that are missing.		
Faithfulness to Constraints	No constraint	Script is unfaithful to the constraint		
	Incoherent steps	Steps that are related to the goal, but are not coherent within the script.		
	Unrelated steps	Steps that are not related to the goal.		

Table 10: The error types and their explanation with examples.

Hyper-parameter Assignment	
Top- p	1.0
Temperature	1.0
Max tokens	512
Presence penalty	0.0
Frequency penalty	0.0
K	2

Table 11: Hyper-parameters for script generation from InstructGPT.




Figure 8: The category distribution of CoScript. The categories are derived from wikiHow

C.2 Domain Examination

As shown in Figure 8, CoScript is derived from wikiHow and encompasses 19 daily life goal categories. These categories cover a wide range of practical topics of everyday life, excluding sensitive and high-stakes topics like medical and legal matters. In addition, we adopt ChatGPT to assess each specific goal in CoScript to further mitigate risks associated with sensitive domains. The instruction is shown in Table 12. Upon examination, ChatGPT identifies that 0.24% of specific goals (132) within the CoScript involve a sensitive domain, such as "relieve head pain with medication" in the *Health* domain. We manually remove these data and sub-

/ Task prompt */*

Given a goal, please carefully evaluate and determine if it falls under sensitive and high-risk domains, specifically within the fields of medical and legal professions.

/ Examples */*

Goal: Treat COVID-19 at Home

Answer: Yes

Goal: Make Strawberry Cake with Microwave

Answer: No

/ Test data */*

Goal: Clean Oil Stains from Clothes with Soap

Answer: **No**

Table 12: The instruction for ChatGPT to identify sensitive and high-risk domains. Generated texts by ChatGPT are **highlighted**.

stitute them with new specific goals (*e.g.*, relieve head pain with meditation) to ensure the safety of our dataset. We encourage researchers and developers to leverage this dataset to build models that accurately understand and respond to user queries on non-sensitive, non-critical topics.

C.3 Qualitative Generations

We randomly select qualitative generations from the CoScript. Table 15, Table 16 and Table 17 show some specific goal generations under different types of constraints. Table 18 shows some scripts generated based on our proposed pipeline.

D Crowd-sourcing Details

Interface Details We conduct human evaluations on Amazon Mechanical Turk. Screenshots of the instructions and annotation interface are shown in Figure 9 and 10.

Step 1: Specific Goal Evaluation We first assess the specific goals generated by InstructGPT with three types of constraints. We ask the turkers to check the specific goal whether inherits the abstract goal and contains a constraint.

Step 2: Script Evaluation In the second step, we show a script of the specific goal with actionable steps. We then ask two questions:

1. *Does the script meet the constraint in the specific goal?* (Yes, No, or Not sure). In our preliminary experiments, we found that the semantic completeness in the scripts generated based on InstructGPT (175B) is acceptable, but faithfulness to the constraints can not be guaranteed. This question assesses this tricky error;
2. *Are the steps in the script correct in achieving the specific goal?* (Yes, No, or Not sure). This question is to assess whether the script can indeed accomplish the given goal. Although we have checked the constraint in the first question, there are still other error types (as shown in Figure 10). Then, we ask the turkers to review the generated scripts. If the scripts cannot achieve the given goal, they must point out the wrong steps and select the error types.

Task Prompt	Create possible Specific Goals according to the Abstract Goal when the Constraint Type is XXX	
	Constraint Type: Modifier	
Example 1: Abstract Goal: Say Goodbye in Different Language Constraint: French Specific Goal: Say Goodbye in French Constraint: English Specific Goal: Say Goodbye in English	Example 2: Abstract Goal: Draw flowers Constraint: Pink Specific Goal: Draw pink flowers Constraint: Blue Specific Goal: Draw blue flowers	Example 3: Abstract Goal: Make hairstyle Constraint: At home Specific Goal: Make hairstyle at home Constraint: At a salon Specific Goal: Make hairstyle at a salon
	Constraint Type: Method	
Example 1: Abstract Goal: Lower blood pressure Constraint: With medication Specific Goal: Lower blood pressure with medication Constraint: With exercises Specific Goal: Lower blood pressure with exercises	Example 2: Abstract Goal: Write a book Constraint: By hand Specific Goal: Write a book by hand Constraint: By typing Specific Goal: Write a book by typing	Example 3: Abstract Goal: Register to vote Constraint: Online Specific Goal: Register to vote online Constraint: Via mail Specific Goal: Register to vote via mail
	Constraint Type: Intent	
Example 1: Abstract Goal: Make a cake Constraint: For a birthday party Specific Goal: Make a cake for a birthday party Constraint: For a wedding Specific Goal: Make a cake for a wedding	Example 2: Abstract Goal: Send an email Constraint: To get a job Specific Goal: Send an email to get a job Constraint: For leave Specific Goal: Send an email for leave	Example 3: Abstract Goal: Buy flowers Constraint: For girlfriend Specific Goal: Buy flowers for girlfriend Constraint: For mother Specific Goal: Buy flowers for mother

Table 13: The handpicked examples for specific goal generation.

Specific Goal: Wear a pencil skirt	
Script (InstructGPT): 1. Pick the right fabric. 2. Choose the right length. 3. Find the right fit. 4. Accessorize appropriately. 5. Wear the right shoes.	Script (Ours): 1. Find a pencil skirt that fits you well. 2. Pair your pencil skirt with a fitted top. 3. Choose the right shoes to complete your look. 4. Accessorize as needed.
Specific Goal: Clean your skin with soap and water	
Script (InstructGPT): 1. Get your skin wet with lukewarm water. 2. Apply a small amount of cleanser to your skin. 3. Rub the cleanser over your skin in gentle, circular motions. 4. Rinse the cleanser off your skin with lukewarm water. 5. Pat your skin dry with a clean, soft cloth.	Script (Ours): 1. Wet your skin with warm water. 2. Apply a dollop of mild soap to your hands. 3. Rub your hands together to create a lather. 4. Apply the lather to your wet skin. 5. Gently massage the lather into your skin in a circular motion. 6. Rinse the soap off your skin with warm water. 7. Pat your skin dry with a clean towel.
Specific Goal: Make stewed fruit with a crockpot	
Script (InstructGPT): 1. Combine all the ingredients in the crockpot. 2. Cook on low for four hours. 3. Serve with ice cream or yogurt.	Script (Ours): 1. Prepare the fruit. 2. Add the sugar. 3. Cook the fruit on low heat in the crockpot. 4. Divide the fruit evenly into four bowls. 5. Make your crumble. 6. Sprinkle the crumble over your fruit and make it.

Table 14: The generation comparison between InstructGPT (175B) with our method. We highlight the constraints in the given examples.

Abstract Goal	Constraints and Specific Goal
Ask a teacher for help	Constraint: Math
	Specific Goal: Ask a math teacher for help
	Constraint: Science
	Specific Goal: Ask a science teacher for help
	Constraint: In school
	Specific Goal: Ask a teacher for help in school
Make pancakes	Constraint: Banana
	Specific Goal: Make banana pancakes
	Constraint: Chocolate chip
	Specific Goal: Make chocolate chip pancakes
	Constraint: Halo 3
	Specific Goal: Download Halo 3 for Xbox 360
Download an xbox 360 game	Constraint: Gears of war
	Specific Goal: Download Gears of War for Xbox 360

Table 15: Qualitative generations for specific goals when the type of constraints is *Modifier*.

Abstract Goal	Constraints and Specific Goal
Prevent kidney stones from recurring	Constraint: By eating a healthy diet
	Specific Goal: Prevent kidney stones from recurring by eating a healthy diet
	Constraint: By taking medication
	Specific Goal: Prevent kidney stones from recurring by taking medication
	Constraint: With a sewing machine
	Specific Goal: Sew Chain Stitch with a sewing machine
Sew chain stitch	Constraint: By hand
	Specific Goal: Sew chain stitch by jand
	Constraint: Formally
Say goodbye in Spanish	Specific Goal: Say goodbye in Spanish formally
	Constraint: Informally
	Specific Goal: Say goodbye in Spanish informally

Table 16: Qualitative generations for specific goals when the type of constraints is *Method*.

Abstract Goal	Constraints and Specific Goal
Use clary sage	Constraint: For aromatherapy
	Specific Goal: Use clary sage for aromatherapy
	Constraint: For skin care
	Specific Goal: Use clary sage for skin care
Make carrot oil	Constraint: For hair care
	Specific Goal: Use clary sage for hair care
	Constraint: For skin
	Specific Goal: Make carrot oil for skin
Acquire abandoned property	Constraint: For cooking
	Specific Goal: Make carrot oil for cooking
	Constraint: For personal use
	Specific Goal: Acquire abandoned property for personal use
	Constraint: For business use
	Specific Goal: Acquire abandoned property for business use

Table 17: Qualitative generations for specific goals when the type of constraints is *Intent*.

Abstract Goal: Link social media accounts on flipboard	
Specific Goal: Link facebook accounts on flipboard	Specific Goal: Link twitter accounts on flipboard
Script: 1. Launch Flipboard. 2. Go to Content. 3. Go to Accounts. 4. Select Facebook.	Script: 1. Launch Flipboard. 2. Go to Content. 3. Go to Accounts. 4. Select Twitter. 5. Log into your Twitter account.
Abstract Goal: Increase your metabolism	
Specific Goal: Increase your metabolism with medication	Specific Goal: Increase your metabolism with diet
Script: 1. Talk to your doctor about prescription options. 2. Follow your doctor's instructions carefully. 3. Stay hydrated	Script: 1. Eat small, frequent meals throughout the day. 2. Choose lean proteins. 3. Add spice to your favorite foods.
Abstract Goal: Use clary sage	
Specific Goal: Use clary sage for aromatherapy	Specific Goal: Use clary sage for massage
Script: 1. Make a clary sage and lavender sachet to keep in your pillowcase. 2. Diffuse clary sage oil to boost your mood and energy levels. 3. Add a few drops of clary sage oil to your bathtub for a relaxing soak	Script: 1. Heat clary sage oil until it is warm. 2. Massage the oil into the muscles and tissues. 3. Leave the oil on for at least 30 minutes. 4. Rinse the oil off with warm water.

Table 18: Qualitative generations for scripts under the different **constraints**.

Instructions (click to expand/collapse)

Thanks for participating in this HIT! **Please spend some time reading this instruction and the example section to better understand our HIT!**

In this hit, you need to complete 20 data labeling tasks. In each task, you will be presented a general goal about everyday activities (such as "make a cake") and a specific goal which inherits the general goal but is more specific and has a reasonable **CONSTRAINT** (such as "make a chocolate cake"). You will answer 3 questions for each task.

Question 1

In Question 1, you need to assess whether the specific goal is reasonable. For example, making a chocolate cake is a reasonable constraint of making a cake, whereas making a lego cake is not reasonable.

Question 2

Then, you will read a script of the specific goal with actionable steps (in the cake' s example, the script is the steps towards making a cake). Question 2 is to check whether the script **MEETS THE CONSTRAINT**. If the specific goal is making a chocolate cake and the script does not mention chocolate, then it does not meet the constraint.

Question 3

In Question 3, you will assess whether the script can indeed accomplish the given goal. If the script can not accomplish the given goal, you need to point out the wrong steps and **SELECT THE ERROR TYPES**. A script for making chocolate cake might mention chocolate, but if its making instructions are wrong, then you need to reflect it in Question 3.

Notes:

- A general goal involves **STEREOTYPICAL ACTIVITIES** such as "make a cake", while a specific goal can be multi-facet **WITH A REASONABLE CONSTRAINT**.
 - For example, a cake can be made for different purposes (for a wedding or a birthday party), with various tools (with a microwave or an oven) or with different ingredients (chocolate or vanilla).
- If you think the specific goal is not reasonable, choose NO in Question 1, but still proceed with Question 2 and 3 pretending that it is reasonable. For example, making a LEGO cake is not reasonable, but you can still assess whether the corresponding script meets the constraint or not. Remember you can always choose "I am not sure" .
- You **SHOULD NOT** ignore grammar and spelling mistakes in the script.
- You can **SEARCH GOOGLE** to help you judge whether the script can achieve the goal, especially if you are not sure about the script.

Example of Error Types

Specific Goal: Make A Vanilla Cake

Script:

1. Gather your ingredients.
2. Preheat the oven to 325 °F (163 °C) and grease and flour a cake pan.
3. Cream the butter and sugar.
4. Add the eggs and vanilla.
5. Stir in the cake flour.
6. Pour the batter into the pan.
7. Bake the cake for 1 hour 15 minutes.

Golden Script

Specific Goal: Make A Vanilla Cake

Script:

1. Preheat the oven to 325 °F (163 °C) and grease and flour a cake pan.
2. Gather your ingredients.
3. Buy your ingredients.
4. Cream the butter and salt.
5. Stir in the cake flour.
6. Have a shower.
7. Pour the batter into the pan.
8. Bake the cake for 1 hour 15 minutes.

Wrong Script

Error Types	Explanation
Wrong order	Steps that are in the wrong order
Repeat steps	Steps that are repeated in the script
Incoherent steps	Steps that are related to the goal, but are not coherent within the script
Unrelated steps	Steps that are not related to the goal
Missing steps	Important steps that are missing
Others: _____ (Describe what you think is wrong)	

I certify that I have read and understand all the instructions.

Figure 9: Instructions for crowd workers on Amazon Mechanical Turk.

Examples (click to expand/collapse)

Question 1

General Goal: Make A Cake

Specific Goal: Make An Apple Cake

Read the given goal. Does the **Specific Goal** inherit the General Goal and contain a **CONSTRAINT (Apple)**?

Select an option
Yes
No
I am not sure

Question 2

Specific Goal: Make An Apple Cake

Script:

1. Gather your ingredients.
2. Preheat the oven to 350 °F (177 °C) and grease and flour the cake pan.
3. In a large bowl, mix the dry ingredients.
4. Cut the apples into small pieces and add them to the bowl.
5. Pour the batter into the cake pan.
6. Bake the cake for 1 hour.
7. Take the cake from the oven and let it cool.
8. Serve

Read the script. Does the script meet the **CONSTRAINT** in the Specific Goal ?

Select an option
Yes
No
I am not sure

Question 3

Read the script. Are the steps in the script correct in achieving the Specific Goal?

Select an option
Yes
No
I am not sure even if I have searched on google

If you choice "NO" , please highlight the wrong steps and select the error types.

Please refer to the instructions for error types.

Figure 10: The examples given in the user interface.

ACL 2023 Responsible NLP Checklist

A For every submission:

- A1. Did you describe the limitations of your work?
Limitations
- A2. Did you discuss any potential risks of your work?
Ethics Statement
- A3. Do the abstract and introduction summarize the paper’s main claims?
Abstract and Section 1
- A4. Have you used AI writing assistants when working on this paper?
Check grammar for the whole paper.

B Did you use or create scientific artifacts?

Section 5.1

- B1. Did you cite the creators of artifacts you used?
Section 4.1
- B2. Did you discuss the license or terms for use and / or distribution of any artifacts?
Section 4.1 and Section 5.1 and Appendix C.1
- B3. Did you discuss if your use of existing artifact(s) was consistent with their intended use, provided that it was specified? For the artifacts you create, do you specify intended use and whether that is compatible with the original access conditions (in particular, derivatives of data accessed for research purposes should not be used outside of research contexts)?
Section 4.1 and Section 5.1 and Appendix C.1
- B4. Did you discuss the steps taken to check whether the data that was collected / used contains any information that names or uniquely identifies individual people or offensive content, and the steps taken to protect / anonymize it?
Ethics Statement
- B5. Did you provide documentation of the artifacts, e.g., coverage of domains, languages, and linguistic phenomena, demographic groups represented, etc.?
Section 5
- B6. Did you report relevant statistics like the number of examples, details of train / test / dev splits, etc. for the data that you used / created? Even for commonly-used benchmark datasets, include the number of examples in train / validation / test splits, as these provide necessary context for a reader to understand experimental results. For example, small differences in accuracy on large test sets may be significant, while on small test sets they may not be.
Section 5

C Did you run computational experiments?

Section 4.4 and Section 6

- C1. Did you report the number of parameters in the models used, the total computational budget (e.g., GPU hours), and computing infrastructure used?
Section 4.4 and Section 6.1 and Appendix C.1

The Responsible NLP Checklist used at ACL 2023 is adopted from NAACL 2022, with the addition of a question on AI writing assistance.

- C2. Did you discuss the experimental setup, including hyperparameter search and best-found hyperparameter values?

No response.

- C3. Did you report descriptive statistics about your results (e.g., error bars around results, summary statistics from sets of experiments), and is it transparent whether you are reporting the max, mean, etc. or just a single run?

No response.

- C4. If you used existing packages (e.g., for preprocessing, for normalization, or for evaluation), did you report the implementation, model, and parameter settings used (e.g., NLTK, Spacy, ROUGE, etc.)?

No response.

D Did you use human annotators (e.g., crowdworkers) or research with human participants?

Appendix D

- D1. Did you report the full text of instructions given to participants, including e.g., screenshots, disclaimers of any risks to participants or annotators, etc.?

Appendix D

- D2. Did you report information about how you recruited (e.g., crowdsourcing platform, students) and paid participants, and discuss if such payment is adequate given the participants' demographic (e.g., country of residence)?

Ethics Statement

- D3. Did you discuss whether and how consent was obtained from people whose data you're using/curating? For example, if you collected data via crowdsourcing, did your instructions to crowdworkers explain how the data would be used?

Appendix D

- D4. Was the data collection protocol approved (or determined exempt) by an ethics review board?

We do not have any human subjects research and use filtered model-generated data as datasets.

- D5. Did you report the basic demographic and geographic characteristics of the annotator population that is the source of the data?

Ethics Statement