

Toward Unified Controllable Text Generation via Regular Expression Instruction

Xin Zheng^{1,3}, Hongyu Lin^{1*}, Xianpei Han^{1,2*}, Le Sun^{1,2}

¹Chinese Information Processing Laboratory ²State Key Laboratory of Computer Science Institute of Software, Chinese Academy of Sciences, Beijing, China

³University of Chinese Academy of Sciences, Beijing, China
{zhengxin2020, hongyu, xianpei, sunle}@iscas.ac.cn

Abstract

Controllable text generation is a fundamental aspect of natural language generation, with numerous methods proposed for different constraint types. However, these approaches often require significant architectural or decoding modifications, making them challenging to apply to additional constraints or resolve different constraint combinations. To address this, our paper introduces Regular Expression Instruction (REI), which utilizes an instruction-based mechanism to fully exploit regular expressions’ advantages to uniformly model diverse constraints. Specifically, our REI supports all popular fine-grained controllable generation constraints, i.e., lexical, positional, and length, as well as their complex combinations, via regular expression-style instructions. Our method only requires fine-tuning on medium-scale language models or few-shot, in-context learning on large language models, and requires no further adjustment when applied to various constraint combinations. Experiments demonstrate that our straightforward approach yields high success rates and adaptability to various constraints while maintaining competitiveness in automatic metrics and outperforming most previous baselines. ¹

1 Introduction

Generating texts according to human requirements has long been a critical challenge in natural language generation (Ziegler et al., 2019; Ouyang et al., 2022). With the emergence of large language models, many tasks in natural language processing can be unified and converted into the formation of *controllable generation* (Prabhumoye et al., 2020). For example, text classification (Apté et al., 1994), cloze test (Devlin et al., 2019), and multiple-choice question answering (Lai et al., 2017) tasks constraint the output text to be exactly one of the given

*Corresponding Authors

¹Our code and data are available at <https://github.com/MrZhengXin/CTG-Regex-Instruction>.

Lexicon & length constraint

Input

<expression> <mask_0> stood(0) <mask_1> field(1)
<mask_2> looking(2) <mask_3> <length=10> </expression>

Output

<expression> The_1 player_2 stood(0)_3 in_4 the_5
field(1)_6 looking(2)_7 at_8 the_9 batter_10 </expression>

Position & lexicon constraint

Input

Stephen was at a party. <expression> <mask_0>
knocked(0) <mask_1> </expression> He checked it but
it was completely broken.

Output

<expression> Stephen knocked(0) over a vase while
drunk. </expression>

Position constraint & alternative ending

Input

My friends all love to go to the club to dance. They think
it’s a lot of fun and always invite. I finally decided to tag
along last Saturday. <expression> <options> <choice_0>
<mask_0> My friends decided to keep inviting me out as
I am so much fun. </choice_0> <choice_1> <mask_1>
The next weekend, I was asked to please stay home.
</choice_1> </options> </expression>

Output

<expression> I danced terribly and broke a friend’s toe.
The next weekend, I was asked to please stay home.
</expression>

Table 1: Input and output of instruction prompt based Regular Expression Instruction (REI). REI can describe various types of complex fine-grain constraints, and here we present three examples. Meta-data instruction label is **colored**, lexicon constraints or correct choice is **bold-faced**, and auxiliary marks for length or lexicon uses gray color.

options; abductive reasoning (Bhagavatula et al., 2020) specifies that the position of the output text is between the previous and future contexts; summarization task (Luhn, 1957) limits the length of output; machine translation (Bar-Hillel, 1960) demands to use the vocabulary of the target language for text generation.

For controllable text generation, typical fine-

grained control tasks include lexicon (Lin et al., 2020), generating position (Shen et al., 2020) and length (Carlsson et al., 2022). Recently, various approaches have been proposed to satisfy these constraints, which can be categorized into three different paradigms: retraining or refactoring the model (Keskar et al., 2019; Zhang et al., 2020; He, 2021; Chan et al., 2021a); tuning on given data (Lester et al., 2021; Stiennon et al., 2020a); manually designed post-processing (Qin et al., 2020, 2022; Meng et al., 2022; Lu et al., 2021, 2022; Wang et al., 2021).

Despite the reasonable performance, current methods on transformer-based language models mainly focus on certain constraints but may not be easily transferred to others, let alone the combination of constraints. For example, Non-Residual Prompting (Carlsson et al., 2022) and A*esque Decoding (Lu et al., 2022) only considered lexical and length constraints, but it cannot arbitrarily specify which position the generated text shall occur; on the other hand, COLD (Qin et al., 2022) can generate text given past and future context, but may not add word inclusion constraint nor restrict the output length. Moreover, these controlling methods assume that we have access to the probability distribution or even gradient of the model, but in the case of large language models where we can only obtain the output token via API, these methods may not be available, and thus black-box controlling techniques need further exploration.

To address the above challenges, we proposed instruction-based Regular Expression Instruction (REI), for universal fine-grained controllable generation. Table 1 present a few examples. Our instruction design is inspired by regular expression, which can easily describe mainstream constraints and their combinations. Following Rosenbaum et al. (2022), we use markup language to construct the expression, hoping that model can better distinguish between meta-data (instructions) and data (actual words). We use two popular paradigms, language model fine-tuning, and large language model few-shot, to teach the model to understand the input constraint expression.

Our method has several advantages. First, our constraint expression supports all typical fine-grained controlling task and is powerful enough to describe composite control specifications. Second, our method can be adapted to various scenarios, such as summarization with length con-

straint, terminology-constrained machine translation, and alternative-ending story infilling. Third, our method is easy to implement and highly transferrable to other models since it requires only fine-tuning on medium-size models and no further modification on large language models, and it does not need access to probability distribution or gradient.

Experiments demonstrate that current state-of-the-art language models can understand our controlling language, achieving high success rate while maintaining high automatic evaluation metric score and surpassing most of the strong previous baselines under various constraints. We hope our work can shed light on future works.

2 Method

2.1 Instruction Design

The controlling language REI follows the style of regular expression due to its expressiveness. Also, it’s easy to evaluate whether the input expression instruction matches the generated text or not. Following Rosenbaum et al. (2022), HTML-like markup language is used, which helps the model learn that they are meaningful meta-data instructions rather than plain symbols, especially when using large language models in-context learning with limited examples and no parameter update. This markup label can also avoid the usage of the escape character.

REI contains several special labels, as shown in Table 1. `<expression>` and `</expression>` mark the beginning and the end of the expression and can be put anywhere in the input text, assuming we only generate according to one expression at a time. `<mask_i>` is equivalent to the regular expression “`.*`” and similar to the mask token in BART (Lewis et al., 2020) and T5 (Raffel et al., 2022), where at its position the model shall generate zero or more tokens. `<options>` and `</options>` is equivalent to the parentheses “`(`” and “`)`” in regular expression, the model shall choose one expression among the group. To make the recognition easier, we use `<choice_i>` and `</choice_i>` to wrap each choice. The regular expression notation of length counts at the character level, but in practice, we want to control the output word length. Therefore, we use the `<length=n>` label to denote the constraint of output word count.

We avoid the shortcoming of T5 (Raffel et al., 2022) span-corruption schema, where the model only generates discontinued spans rather than full

Task	Input with Control Expression
α NLG	O_1 <expression> <mask_0> </expression> O_2
α NLG+length	O_1 <expression> <mask_0> <length= l > </expression> O_2
α NLI	O_1 <expression> <options> <choice_0> H_1 </choice_0> <choice_1> H_2 </choice_1> </options> </expression> O_2
CommonGen	<expression> <mask_0> $c_0(0)$ <mask_1> $c_1(1)$ <mask_2> $c_2(2)$ <mask_3> </expression>
CommonGen+length	<expression> <mask_0> $c_0(0)$ <mask_1> $c_1(1)$ <mask_2> $c_2(2)$ <mask_3> <length= l > </expression>

(a) Fine-tune Task

Task	Input with Control Expression
α NLG+lexicon	O_1 <expression> <mask_0> $w(0)$ <mask_1> </expression> O_2
α NLG+length+lexicon	O_1 <expression> <mask_0> $w(0)$ <mask_1> <length= l > </expression> O_2
StoryCompletion+infill	$S_1 S_2 S_3$ <expression> <mask_0> <options> <choice_0> E_1 </choice_0> <choice_1> E_2 </choice_1> </options> </expression>
Gigaword+length	[Text]\n Summarize the aforementioned text in a single phrase.\n <expression> <mask_0> <length= l > </expression>
Wiktionary/ITAE	Translate from English to German:\n\n English: [Text] \n German: <expression> <mask_0> $t_0(0)$ <mask_1> $t_1(1)$ <mask_2> </expression>

(b) Transfer Task

Table 2: Constraint expression of each task. We fine-tune on tasks and variations listed in Table 2a, and additionally evaluate the unseen tasks listed in Table 2b. Notice that for few-shot learning, all the tasks are not trained before.

natural sentences (Lester et al., 2021). On the other hand, we also overcome the redundancy of BART denoising schema (He, 2021), where the whole input is generated again, since we only generate the realized expression. Moreover, beyond fill-in-the-blank, we introduce choice-making, which further enriches the expressiveness of our controlling language.

2.2 Training

Fine-tuning We could automatically construct the training data from the corpus and conduct self-supervised learning. Alternatively, we could also directly convert the input of existing supervised datasets into the form of our controlling language, and use them to fine-tune state-of-the-art models such as FLAN-T5 (Chung et al., 2022). The input format is shown in Table 2a.

We include α NLG (Bhagavatula et al., 2020) and CommonGen (Lin et al., 2020), two English controllable generation datasets of position and lexicon constraint. In α NLG, given the past observation O_1 and the future observation O_2 , the goal is to generate a hypothesis h that could follow O_1 and trigger O_2 . The regular expression of the constraint is “.*” since no lexicon constraint is required. In CommonGen, given a set of k concepts $C = \{c_0, c_1, \dots, c_{k-1}\}$, the output text shall include those concepts and at the same time be consistent with common sense. While in the original setting, the appearance order of concepts and their

word sense change is not provided, and the model shall make these decisions, here in our controlling language, the exact word and order must be given. Otherwise, we cannot construct the corresponding expression. So, we preprocess the original instances and recover the order and word sense of the concepts by the reference text. To help the model generate the concepts sequentially and track how many concepts it has already used, we append the serial number label (i) to every concept c_i on both the input and output sides and remove the labels from the output generation once completed. The regular expression of the constraint is “. * c_0 . * c_1 * c_{k-1} . *”.

We also leverage these two datasets to teach the model to control the output length by simply adding the length label with the ground truth length. To better track how many words the model itself has already generated, we append the length number label $_i$ to every word w_i ; for example, the sentence “Stephen knocked over a vase while drunk.” becomes “Stephen_0 knocked_1 over_2 a_3 vase_4 while_5 drunk_6”. Similarly, we remove the length number labels after completion.

Finally, we need to teach the model about choosing grammar. We use α NLI (Bhagavatula et al., 2020) dataset, the task of which is to determine whether H_1 or H_2 is the more plausible hypothesis given the past and future observations O_1 and O_2 , and the constraint of the regular expression is “($H_1 | H_2$)”.

In-context Learning For large language models like GPT-3.5 (Brown et al., 2020), where typically access is typically provided via API, we may not apply many traditional controllable generation techniques. However, we can leverage its ability of in-context learning to conduct fine-grain constraint generation. More specifically, we leverage the ability to discover and imitate the repeated pattern (Madaan and Yazdanbakhsh, 2022; Min et al., 2022), which is desirable in our case, since unlike other natural language understanding tasks, the specific fine-grain constraint is a well-defined simple pattern that could be easily discoverable and imitable.

Given the input with control expression, we can select k instances with the same expression structure as the instruction prompt and send it to the large language model together with input. Naturally, when evaluating the test set, we can select examples from the training set or validation set, or other instances of the test set when they are not available. Consistently, we use the same input and output format described before, which saves extra efforts on prompt engineering. In addition, we simply use the popular json format “ {”input”: [INPUT], ”output”: [OUTPUT]} ” for each demonstrating instances, and naturally separate them with “\n”. By using json, we can further avoid the need for escape character if the input text happens to contain metadata like "Input" or "\n".

2.3 Inference

We use rejection sampling to generate output text that is matched by the control expression. Verifying the output is simple, since we could convert the control expression into regular expression and check the validity. Additionally, if the expression contains length constraint label, we count and compare the number of words in the output text. We try at most k times to avoid infinite loop and save costs if we use large language model API. When using medium or small size language model, to increase the generation quality, we can perform beam search first and see if it can generate a valid result at the first try.

2.4 Recursive Decoding

Different choice might affect the generated text. For example, consider the case “ $S_1 S_2 S_3 . *(E_1 | E_2)$ ”, which gives the first three sentence and two alternative endings and the goal is to choose the correct ending while infill

the fourth sentence at the same, which is not included in our fine-tuning data. Instead of directly jumping to the answer with possibly insufficient computation, we could also let the model “think step by step (Kojima et al., 2022)”. We can solve each choice expression first, then compare the complete choices “ $(S_4 E_1 | S_4' E_2)$ ”. The generalized decoding procedure is presented at Algorithm 1, which assumes that each options is independent with each other and greedily solve them from left to right. We leave the evaluation of expression with multiple consecutive options (Lu et al., 2022) for future work.

3 Experiment

3.1 Setup

We conduct experiments on 2 Nvidia A100 GPUs, with about 10 total GPU hours locally. For medium-size language model, we use FLAN-T5-xl (Chung et al., 2022) with Apache 2.0 license, which has 3B parameters and is fine-tuned on many natural language understanding and generation tasks. We use Huggingface Transformers library (Wolf et al., 2020) with Apache-2.0 license for fine-tuning and evaluation. We trained the model for 3 epochs, with a batch size of 16 and learning rate of $3e-5$. We set beam size to 4 for beam search and p to 0.95 for top- p sampling. We generate at most $k = 512$ samples if we do not obtain any valid outcome.

For large language model, we use GPT-3 (Brown et al., 2020) text-davinci-003 version via OpenAI API, and the 175B model is calibrated with Reinforcement Learning from Human Feedback (Stiennon et al., 2020b). We feed 8 in-domain examples as the prompt, set the temperature to 0.7, and retry at most $k = 8$ times if the result is not valid. All results are from “single” run.

3.2 Lexicon Constraint

3.2.1 Lexicon Constraint Only

Setting We evaluate our method on the devset of CommonGen (Lin et al., 2020), as the reference text of the test set is not publicly available. As mentioned in 2.2, we feed the model with oracle concept order and word sense. For automatic metrics we use BLEU-4 (Papineni et al., 2002), CIDEr (Vedantam et al., 2015), SPICE (Anderson et al., 2016) and Coverage (Cov.), which is the average ratio of input concepts that are present in lemmatized outputs.

Method	BLEU	CIDEr	SPICE	Cov.
BART (Lin et al., 2020)	31.83	13.96	28.00	97.35
T5-Large (Lin et al., 2020)	31.96	15.13	28.86	95.29
Neurologic (Lu et al., 2021)	28.10	15.50	30.80	98.50
NADO (Meng et al., 2022)	30.80	-	-	97.10
NRP (Carlsson et al., 2022)	-	-	-	95.10
NLI+GPT-3.5, 8 shot, oracle	38.89	18.60	31.51	98.93
REI+GPT-3.5, 8 shot, oracle	28.64	15.15	29.49	98.60
REI+FLAN-T5-xl, oracle	36.78	18.34	33.56	100.0

(a) Lexicon constraint

Method	BLEU	CIDEr	SPICE	SuR.
NLI+GPT-3.5, 8 shot	40.48	19.72	31.78	35.95
REI+GPT-3.5, 8 shot	19.53	11.54	22.35	67.43
REI+FLAN-T5-xl	30.95	17.50	32.37	99.90

(b) Lexicon & length constraint

Table 3: Results on devset of CommonGen. The best models are **bold** within each metric.

Results We compare the performance of our method with other baselines, including fine-tuning methods BART (Lin et al., 2020) and T5-Large (Lin et al., 2020), auxiliary guiding model method NADO (Meng et al., 2022), prompting method NRP (Carlsson et al., 2022), and 8-shot pure natural language instruction (NLI) on GPT-3.5, which is shown at Table 3a.

Given only 8 examples with a clear connection between input and output, GPT-3.5 still shows competitive performance in terms of text automatic metrics, and achieves high concept coverage, surpassing all the previous baselines. Compared with natural language instruction, the success rate is very close. And with more supervised data to modify the model’s parameter, FLAN-T5-xl performs significantly better than GPT-3.5 and other previous baselines in all metrics and successfully satisfies all lexicon constraints.

3.2.2 Lexicon & Length Constraint

As described in Section 2.2, we slightly modify the devset of CommonGen to introduce the additional length constraint and evaluate GPT-3.5 and FLAN-T5. For metric, we replace Coverage (Cov.) with Success Rate (SuR.), which is the average percentage of output that matches the input expression. In a composite task, the performance of GPT-3.5 downgrades dramatically and struggles to generate valid output, indicating that multi-concept inclusion and length control at the same time is challenging, especially for few-shot in-context learning. Yet, REI still outperforms NLI in terms of success rate, and the "high" n-gram metrics might also

indicate the poor instruction following ability in terms of challenging fine-grain constraints, which is consistent with the finding of Zhou et al. (2023). FLAN-T5 only has a minor drop in performance and still maintains a high success rate since it has trained on this composite constraint.

3.3 Position Constraint

3.3.1 Position constraint only

Setting We evaluate our method on the testset of α NLG (Bhagavatula et al., 2020). The automatic metrics include BLEU-4 (Papineni et al., 2002), ROUGE-L (Lin, 2004) and BERTScore (Zhang* et al., 2020). We do not report Success Rate since it’s always 100%.

Results As presented in Table 4a, we compare our method with two unsupervised baselines DeLorean (Qin et al., 2020) and COLD (Qin et al., 2022), non-autoregressive Diffusion-LM (Li et al., 2022) and two fine-tuned methods on 11B T5 (Khashabi et al., 2021), 20B UL2 (Tay et al., 2022) and 8-shot NLI on GPT-3.5.

With few-shot learning, GPT-3.5 outperforms two unsupervised baselines and Diffusion-LM, demonstrating its strong in-context learning ability given only a few infilling examples. Since it’s a relatively simple constraint, the performance between REI and NLI is very close. With our careful instruction prompt design and adequate fine-tuning, 3B FLAN-T5 shows stronger performance than 11B T5, and remains competitive compared to 20B UL2.

Method	BLEU	ROUGE	BERT
Qin et al. (2020)	1.38	18.94	42.86
Qin et al. (2022)	1.79	19.50	42.67
Li et al. (2022)	7.10	28.30	89.00
Khashabi et al. (2021)	19.47	44.60	92.87
Tay et al. (2022)	24.34	49.30	93.51
NLI+GPT-3.5, 8 shot	13.62	36.38	91.05
REI+GPT-3.5, 8 shot	13.01	37.29	91.27
REI+FLAN-T5-xl	25.44	48.45	93.28

(a) Position constraint

Model	BLEU	ROUGE	SuR.
NLI+GPT-3.5, 8 shot	9.9	32.93	42.09
REI+GPT-3.5, 8 shot	10.63	34.87	96.80
REI+FLAN-T5-xl	19.92	46.17	100.0

(b) Position & length constraint

Model	BLEU	ROUGE	SuR.
NLI+GPT-3.5, 8 shot	14.76	42.04	99.01
REI+GPT-3.5, 8 shot	18.59	44.67	99.44
REI+FLAN-T5-xl	23.56	48.81	99.78

(c) Position & lexicon constraint

Model	BLEU	ROUGE	SuR.
NLI+GPT-3.5, 8 shot	19.14	43.67	28.00
REI+GPT-3.5, 8 shot	17.45	43.90	94.02
REI+FLAN-T5-xl	21.99	49.17	99.69

(d) Position & length & lexicon constraint

Table 4: Result on test of α NLG.

3.3.2 Position & Length Constraint

As mentioned in Section 2.2, we slightly modify the α NLG test set to add the length constraint. We change the BERTScore metric to SuccessRate (SuR.). Table 4b shows the results. GPT-3.5 manages to imitate both position and length constraints, showing relatively high success rate, while under NLI, it performs badly. But with full-scale supervised learning, FLAN-T5 can robustly generate valid output on the test set 100% of the time. Also, in terms of automatic metrics, the output of both models does not downgrade dramatically.

3.3.3 Position & Lexicon Constraint

We can also modify the α NLG test set to add lexicon constraint, setting the keyword to be the first verb on the reference text. The input format is shown in Table 2b, and Table 4c shows the results. For GPT-3.5, it still is very likely to generate valid output nearly all of the time, and the automatic metrics enjoy improvement compared with the results of no lexicon constraint, since the additional gold words are provided, and the verb constraint limits the vast scope of possible hypothesis space. Also,

REI is slightly better than NLI. For FLAN-T5, although it has been trained on position constraint or lexicon constraint separately, it has not seen the combination, and yet still demonstrates strong performance.

3.3.4 Position & Lexicon & Length Constraint

We can further combine all conditions together, adding both length and lexicon constraints on the test set of α NLG. The input format is presented in Table 2b, and Table 4d shows the results. Compositional constraints challenge few-shot GPT-3.5, as it’s more difficult to generate output that matches all three requirements, and the success rate drops slightly. Interestingly, NLI got a very low success rate. But fully-trained FLAN-T5 exhibits robust transfer ability, as the simultaneous three constraints are not included in training data, but FLAN-T5 still manages to achieve close to 100% success rate.

3.3.5 Position Constraint & Alternative Endings

On the test set of Story Cloze Test (Mostafazadeh et al., 2016), which is to choose between the right ending and the wrong one given the four-sentence context, we additionally mask the fourth sentence and require the model to infill the missing sentence while determining the correct ending. The input format is shown in Table 2b, and the result is shown in Table 6. We change the Success Rate (SuR.) metric to Accuracy (Acc.), since choosing either ending is valid. For GPT-3.5, we directly construct promoting examples with the initial input and final output, and surprisingly find that GPT-3.5 handles the composite constraint quite well, and chooses the right ending with not bad accuracy. Also, REI comes close to NLI in performance. For FLAN-T5-xl, we use the recursive decoding (Section 2.4, and it shows moderate performance, with lower accuracy but higher BLEU / ROUGE compared with GPT-3.5.

3.4 Summarization with length constraint

REI can also easily support abstractive summarization with desired length (Kikuchi et al., 2016; Fan et al., 2018), as long as the base model has been trained on the summarization task, which is the case in our choosing models FLAN-T5 (Chung et al., 2022) and GPT-3.5 (Ouyang et al., 2022). We choose to evaluate on the test set of English headline generation dataset Gigaword (Graff et al.,

Method	Wiktionary		IATE	
	Term%	BLEU	Term%	BLEU
Constraint decoding (Dinu et al., 2019)	99.50	25.80	82.00	25.30
Train-by-replace (Dinu et al., 2019)	93.40	26.30	94.50	26.00
RePP (Sun et al., 2022)	93.67	30.52	95.41	29.38
TADA (Ailem et al., 2021)	96.84	26.73	98.02	27.11
EDITOR (Xu and Carpuat, 2021)	99.8	29.30	100.0	28.90
Levenshtein Transformer (Susanto et al., 2020)	100.0	31.20	100.0	30.13
NLI+GPT-3.5, 8-shot	99.03	37.62	98.07	32.22
REI+GPT-3.5, 8-shot	99.52	34.88	99.45	35.25

Table 5: Results on Wiktionary and IATE.

Method	BLEU	ROUGE	Acc.
NLI+GPT-3.5, 8 shot	3.83	21.27	88.99
REI+GPT-3.5, 8 shot	3.77	20.56	88.72
REI+FLAN-T5-xl	3.87	20.9	84.61

Table 6: Results on Story Cloze Test with positional constraint.

Method	ROUGE	SuR.
SEQ (Baziotis et al., 2019)	22.68	-
TED (Yang et al., 2020)	22.83	-
NLI+GPT-3.5, 8 shot	24.62	28.87
REI+GPT-3.5, 8 shot	25.46	79.51
REI+FLAN-T5-xl	28.49	100.0

Table 7: Results on the test set of Gigaword.

2003), due to its short input and output length. Also, Gigaword is not included in the training set of FLAN-T5 or GPT-3.5. The input format is written in Table 2b. We use ROUGE-L (Lin, 2004) and Success Rate (SuR.) for metrics.

We compare our methods with two unsupervised unconstrained baselines SEQ (Baziotis et al., 2019) and TED (Yang et al., 2020), and the results are shown in Table 7. Both GPT-3.5 and FLAN-T5 exceed the two baselines in ROUGE-L score, showing relatively good text quality. Since the summarization task constrains more on the semantic of output compared with pure lexicon constraint (CommonGen) or position constraint (α NLG), satisfying length constraint might be more difficult, and GPT-3.5 shows a relatively lower success rate, but NLI has the worst success rate. But nevertheless, FLAN-T5 still achieves 100% success rate. Notice that with limited REI training tasks, the model can still generalize to new tasks with the specific format, demonstrating the robust transfer ability under supervised learning.

3.5 Terminology-constrained machine translation

We can also apply REI to machine translation with terminology constraint (Dinu et al., 2019), which is to ensure the given terminologies $T = (t_0, t_1, \dots)$ are used in translation. We only test GPT-3.5 here, due to its superiority in multi-language understanding, while the majority of output language during pre-training, multi-task learning, and fine-tuning is English. We evaluate on the test set of Wiktionary and IATE (Dinu et al., 2019), two English-German translation dataset, using BLEU-4 (Papineni et al., 2002) and Terminology Coverage (Term) for metrics.

We compare our method with several strong baselines, including Constraint decoding (Dinu et al., 2019), Train-by-replace (Dinu et al., 2019), RePP (Sun et al., 2022), TADA (Ailem et al., 2021), EDITOR (Xu and Carpuat, 2021), Levenshtein Transformer (Susanto et al., 2020), and 8-shot NLI on GPT-3.5. Due to its vast parameters, GPT-3.5 outperforms all other baselines in terms of BLEU score. Also, GPT-3.5 achieves near 100% terminology coverage rate, which is close to the existing upper limit. Finally, REI has a slightly higher term coverage than NLI.

3.6 Qualitative Results

Table 8 shows the samples of lexicon & length constraints (Section 3.2.2), position & lexicon & length constraints (Section 3.3.4), position constraint with alternative ending (Section 3.3.5), summarization with length constraint (Section 3.4) and translation with terminology constraint (Section 3.5). Both FLAN-T5 and GPT-3.5 generate valid and fluent sentences. GPT-3.5 also uses more vivid or human-like words like “antihistamines” or the abbreviation “FIA”, probably due to its large-scale model size and training corpus.

CommonGen+length	<code><expression> <mask_0> dance(0) <mask_1> performed(1) <mask_2> stage(2) <mask_3> wearing(3) <mask_4> costumes(4) <mask_5> <length=11> </expression></code>
FLAN-T5-xl	A_1 dance(0) _2 is_3 performed(1) _4 on_5 a_6 stage(2) _7 by_8 people_9 wearing(3) _10 costumes(4) _11
GPT-3.5, 8 shot	A_1 traditional_2 dance(0) _3 is_4 performed(1) _5 on_6 the_7 stage(2) _8 wearing(3) _9 colorful_10 costumes(4) _11
α NLG+length+lexicon	Jim was not confident in his home repair skills. <code><expression> <mask_0> attended(0) <mask_1> <length=9> </expression></code> Jim was so excited to learn a new skill.
FLAN-T5-xl	Jim_1 bought_2 new_3 gloves_4 and_5 attended(0) _6 a_7 home_8 repair_9
GPT-3.5, 8 shot	Jim_1 attended(0) _2 a_3 home_4 repair_5 workshop_6 to_7 gain_8 confidence_9
StoryCompletion+infill	I tried going to the park the other day. The weather seemed nice enough for a walk. Within minutes of getting there I started sneezing. <code><expression> <options> <choice_0> <mask_0> My allergies were too bad and I had to go back home. </choice_0> <choice_1> <mask_1></code> It reminded me of how much I loved spring flowers. <code></choice_1> </options> </expression></code>
FLAN-T5-xl	There were a lot of people at the park. My allergies were too bad and I had to go back home.
GPT-3.5, 8 shot	I realized I had forgotten the antihistamines at home. My allergies were too bad and I had to go back home.
Gigaword+length	japan 's toyota team europe were banned from the world rally championship for one year here on friday in a crushing ruling by the world council of the international automobile federation.\n Summarize the aforementioned text in a single phrase.\n <code><expression> <mask_0> <length=6> </expression></code>
FLAN-T5-xl	toyota_1 team_2 europe_3 banned_4 from_5 rallying_6
GPT-3.5, 8 shot	toyota_1 team_2 europe_3 banned_4 by_5 fia_6
Wiktionary	Translate from English to German:\n\n English: Jennifer Aniston need not always be perfect or successful. \n German: <code><expression> <mask_0> erfolgreich(0) <mask_1> </expression></code>
GPT-3.5, 8 shot	Jennifer Aniston muss nicht immer perfekt oder erfolgreich(0) sein.

Table 8: Qualitative examples of various constraints by fine-tuned FLAN-T5-xl and few-shot GPT-3.5.

4 Related Work

Tasks of Controllable Text Generation Controllable text generation refers to the tasks that generate text according to the controlling signals (Prabhume et al., 2020). Typically, the output can be constrained at three levels from coarse to fine: (Zhang et al., 2022) semantic, structural and lexical. At semantic level, the signals include topic (Tang et al., 2019), sentiment (Logeswaran et al., 2018), format (Li et al., 2020), toxicity (Krause et al., 2021) and other abstract attribute. At the structural level, the constraints include key-value data table (Novikova et al., 2017), syntax tree, and parts-of-speech (Li et al., 2022). At lexical level, then controlling elements include keyword (Lin et al., 2020), generating position (Shen et al., 2020) and length (Carlsson et al., 2022).

Methods of Controllable Text Generation Current approach for controllable text generation can be summarized as three main categories (Zhang et al., 2022): retraining or refactoring the model, e.g. CTRL (Keskar et al., 2019), POINTER (Zhang et al., 2020), CMDP (Chan et al., 2021b), Constrained BART (He, 2021), CoCon (Chan et al., 2021a), PlanGen (Su et al., 2021) and InstructCTG (Zhou et al., 2023); tuning on given data, including model fine-tuning, Prompt Tuning (Lester et al.,

2021) and RL-Fine Tuning (Stiennon et al., 2020a); and post-processing, which can either design specific decoding strategy, e.g. Constrained Beam Search (Anderson et al., 2017), DeLorean (Qin et al., 2020), COLD (Qin et al., 2022), NeuroLogic (Lu et al., 2021); or using auxiliary guiding model, e.g. PPLM (Anderson et al., 2017), GeDI (Krause et al., 2021), FUDGE (Yang and Klein, 2021), CTRLsum (He et al., 2022), Plug-and-Play Content Planning (Liu et al., 2022), NADO (Meng et al., 2022), and MACSum (Zhang et al., 2023) .

5 Conclusion

We proposed Regular Expression Instruction (REI), a novel instruction-based method that unifies fine-grain lexical-level constrained text generation. Our method is highly adaptable, fitting either language model fine-tuning or large language model in-context learning. Our controlling language can also easily be applied to other related tasks, including story completion while infilling, summarization with length constraint, and machine translation with terminology constraint. Experiments show that our method has a high success rate and outperforms most of the previous strong baselines, demonstrating its effectiveness despite the simplicity. We leave the evaluation and improvement of more complex constraints for future works.

Limitations

Our proposed Regular Expression Instruction is serialized and cannot describe a set of keyword constraints where the appearing order is arbitrary, but only a list of keywords with determined order. Future work is needed to exceed the limit, either by approximating the word order or by repeated random sampling. Also, to obtain valid results we use reject sampling, which might need many repeated trials, thus reducing the efficiency and downgrading the speed. More efficient mechanisms with less retry are worth investigating. Additionally, under the current trends of the instruction following, more sophisticated prompts under 0-shot is worth investigating.

Ethics Statement

This work involves no sensitive data and uses several public-available datasets. This work discusses controllable text generation, which aims for better usage of the black-box language model and may better reduce the problematic biases. We notice that the method proposed in this work can be used to generate disinformation or harmful content directly via controlling language, but the malicious usage can be further avoided by filtering out improper control input and stopping harmful content generation.

References

- Melissa Ailem, Jingshu Liu, and Raheel Qader. 2021. [Encouraging neural machine translation to satisfy terminology constraints](#). In *Findings of the Association for Computational Linguistics: ACL-IJCNLP 2021*, pages 1450–1455, Online. Association for Computational Linguistics.
- Peter Anderson, Basura Fernando, Mark Johnson, and Stephen Gould. 2016. [Spice: Semantic propositional image caption evaluation](#). In *Computer Vision – ECCV 2016*, pages 382–398, Cham. Springer International Publishing.
- Peter Anderson, Basura Fernando, Mark Johnson, and Stephen Gould. 2017. [Guided open vocabulary image captioning with constrained beam search](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 936–945, Copenhagen, Denmark. Association for Computational Linguistics.
- Chidanand Apté, Fred Damerau, and Sholom M. Weiss. 1994. [Automated learning of decision rules for text categorization](#). *ACM Trans. Inf. Syst.*, 12(3):233–251.
- Yehoshua Bar-Hillel. 1960. [The present status of automatic translation of languages](#). This article was prepared with the sponsorship of the information systems branch, office of naval research, under contract nr 049130. reproduction as a whole or in part for the purposes of the u. s. government is permitted. volume 1 of *Advances in Computers*, pages 91–163. Elsevier.
- Christos Baziotis, Ion Androutsopoulos, Ioannis Konstas, and Alexandros Potamianos. 2019. [SEQ³: Differentiable sequence-to-sequence-to-sequence autoencoder for unsupervised abstractive sentence compression](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 673–681, Minneapolis, Minnesota. Association for Computational Linguistics.
- Chandra Bhagavatula, Ronan Le Bras, Chaitanya Malaviya, Keisuke Sakaguchi, Ari Holtzman, Hannah Rashkin, Doug Downey, Wen tau Yih, and Yejin Choi. 2020. [Abductive commonsense reasoning](#). In *ICLR*.
- 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.
- Fredrik Carlsson, Joey Öhman, Fangyu Liu, Severine Verlinden, Joakim Nivre, and Magnus Sahlgren. 2022. [Fine-grained controllable text generation using non-residual prompting](#). In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6837–6857, Dublin, Ireland. Association for Computational Linguistics.
- Alvin Chan, Yew-Soon Ong, Bill Pung, Aston Zhang, and Jie Fu. 2021a. [Cocon: A self-supervised approach for controlled text generation](#). In *International Conference on Learning Representations*.
- Hou Pong Chan, Lu Wang, and Irwin King. 2021b. [Controllable summarization with constrained Markov decision process](#). *Transactions of the Association for Computational Linguistics*, 9:1213–1232.
- Hyung Won Chung, Le Hou, Shayne Longpre, Barret Zoph, Yi Tay, William Fedus, Eric Li, Xuezhi Wang, Mostafa Dehghani, Siddhartha Brahma, Albert Webson, Shixiang Shane Gu, Zhuyun Dai, Mirac Suzgun, Xinyun Chen, Aakanksha Chowdhery, Sharan

- Narang, Gaurav Mishra, Adams Yu, Vincent Y. Zhao, Yanping Huang, Andrew M. Dai, Hongkun Yu, Slav Petrov, Ed H. Chi, Jeff Dean, Jacob Devlin, Adam Roberts, Denny Zhou, Quoc V. Le, and Jason Wei. 2022. [Scaling instruction-finetuned language models](#). *CoRR*, abs/2210.11416.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Georgiana Dinu, Prashant Mathur, Marcello Federico, and Yaser Al-Onaizan. 2019. [Training neural machine translation to apply terminology constraints](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 3063–3068, Florence, Italy. Association for Computational Linguistics.
- Angela Fan, David Grangier, and Michael Auli. 2018. [Controllable abstractive summarization](#). In *Proceedings of the 2nd Workshop on Neural Machine Translation and Generation*, pages 45–54, Melbourne, Australia. Association for Computational Linguistics.
- David Graff, Junbo Kong, Ke Chen, and Kazuaki Maeda. 2003. English gigaword. *Linguistic Data Consortium, Philadelphia*, 4(1):34.
- Junxian He, Wojciech Kryscinski, Bryan McCann, Nazneen Rajani, and Caiming Xiong. 2022. [CTRL-sum: Towards generic controllable text summarization](#). In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 5879–5915, Abu Dhabi, United Arab Emirates. Association for Computational Linguistics.
- Xingwei He. 2021. [Parallel refinements for lexically constrained text generation with BART](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 8653–8666, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Nitish Shirish Keskar, Bryan McCann, Lav R. Varshney, Caiming Xiong, and Richard Socher. 2019. [CTRL: A conditional transformer language model for controllable generation](#). *CoRR*, abs/1909.05858.
- Daniel Khashabi, Gabriel Stanovsky, Jonathan Bragg, Nicholas Lourie, Jungo Kasai, Yejin Choi, Noah A Smith, and Daniel S Weld. 2021. Genie: A leaderboard for human-in-the-loop evaluation of text generation. *arXiv preprint arXiv:2101.06561*.
- Yuta Kikuchi, Graham Neubig, Ryohei Sasano, Hiroya Takamura, and Manabu Okumura. 2016. [Controlling output length in neural encoder-decoders](#). In *Proceedings of the 2016 Conference on Empirical Methods in Natural Language Processing*, pages 1328–1338, Austin, Texas. 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 *Advances in Neural Information Processing Systems*.
- Ben Krause, Akhilesh Deepak Gotmare, Bryan McCann, Nitish Shirish Keskar, Shafiq Joty, Richard Socher, and Nazneen Fatema Rajani. 2021. [GeDi: Generative discriminator guided sequence generation](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 4929–4952, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Guokun Lai, Qizhe Xie, Hanxiao Liu, Yiming Yang, and Eduard Hovy. 2017. [RACE: Large-scale Reading comprehension dataset from examinations](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 785–794, Copenhagen, Denmark. Association for Computational Linguistics.
- Brian Lester, Rami Al-Rfou, and Noah Constant. 2021. [The power of scale for parameter-efficient prompt tuning](#). In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 3045–3059, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Mike Lewis, Yinhan Liu, Naman Goyal, Marjan Ghazvininejad, Abdelrahman Mohamed, Omer Levy, Veselin Stoyanov, and Luke Zettlemoyer. 2020. [BART: Denoising sequence-to-sequence pre-training for natural language generation, translation, and comprehension](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 7871–7880, Online. Association for Computational Linguistics.
- Piji Li, Haisong Zhang, Xiaojiang Liu, and Shuming Shi. 2020. [Rigid formats controlled text generation](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 742–751, Online. Association for Computational Linguistics.
- Xiang Lisa Li, John Thickstun, Ishaan Gulrajani, Percy Liang, and Tatsunori Hashimoto. 2022. [Diffusion-LM improves controllable text generation](#). In *Advances in Neural Information Processing Systems*.
- Bill Yuchen Lin, Wangchunshu Zhou, Ming Shen, Pei Zhou, Chandra Bhagavatula, Yejin Choi, and Xiang Ren. 2020. [CommonGen: A constrained text generation challenge for generative commonsense reasoning](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1823–1840, Online. Association for Computational Linguistics.

- 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.
- Yinhong Liu, Yixuan Su, Ehsan Shareghi, and Nigel Collier. 2022. [Plug-and-play recipe generation with content planning](#). In *Proceedings of the 2nd Workshop on Natural Language Generation, Evaluation, and Metrics (GEM)*, pages 223–234, Abu Dhabi, United Arab Emirates (Hybrid). Association for Computational Linguistics.
- Lajanugen Logeswaran, Honglak Lee, and Samy Bengio. 2018. Content preserving text generation with attribute controls. In *Proceedings of the 32nd International Conference on Neural Information Processing Systems, NIPS’18*, page 5108–5118, Red Hook, NY, USA. Curran Associates Inc.
- Ximing Lu, Sean Welleck, Peter West, Liwei Jiang, Jungo Kasai, Daniel Khashabi, Ronan Le Bras, Lianhui Qin, Youngjae Yu, Rowan Zellers, Noah A. Smith, and Yejin Choi. 2022. [NeuroLogic a*esque decoding: Constrained text generation with lookahead heuristics](#). In *Proceedings of the 2022 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 780–799, Seattle, United States. Association for Computational Linguistics.
- Ximing Lu, Peter West, Rowan Zellers, Ronan Le Bras, Chandra Bhagavatula, and Yejin Choi. 2021. [NeuroLogic decoding: \(un\)supervised neural text generation with predicate logic constraints](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 4288–4299, Online. Association for Computational Linguistics.
- H. P. Luhn. 1957. [A statistical approach to mechanized encoding and searching of literary information](#). *IBM Journal of Research and Development*, 1(4):309–317.
- Aman Madaan and Amir Yazdanbakhsh. 2022. [Text and patterns: For effective chain of thought, it takes two to tango](#). *CoRR*, abs/2209.07686.
- Tao Meng, Sidi Lu, Nanyun Peng, and Kai-Wei Chang. 2022. [Controllable text generation with neurally-decomposed oracle](#). In *Advances in Neural Information Processing Systems*.
- 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?](#) *CoRR*, abs/2202.12837.
- Nasrin Mostafazadeh, Nathanael Chambers, Xiaodong He, Devi Parikh, Dhruv Batra, Lucy Vanderwende, Pushmeet Kohli, and James Allen. 2016. [A corpus and cloze evaluation for deeper understanding of commonsense stories](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 839–849, San Diego, California. Association for Computational Linguistics.
- Jekaterina Novikova, Ondřej Dušek, and Verena Rieser. 2017. [The E2E dataset: New challenges for end-to-end generation](#). In *Proceedings of the 18th Annual SIGdial Meeting on Discourse and Dialogue*, pages 201–206, Saarbrücken, Germany. 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, John Schulman, Jacob Hilton, Fraser Kelton, Luke Miller, Maddie Simens, Amanda Askell, Peter Welinder, Paul F. Christiano, Jan Leike, and Ryan Lowe. 2022. [Training language models to follow instructions with human feedback](#). *CoRR*, abs/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.
- Shrimai Prabhumoye, Alan W Black, and Ruslan Salakhutdinov. 2020. [Exploring controllable text generation techniques](#). In *Proceedings of the 28th International Conference on Computational Linguistics*, pages 1–14, Barcelona, Spain (Online). International Committee on Computational Linguistics.
- Lianhui Qin, Vered Shwartz, Peter West, Chandra Bhagavatula, Jena D. Hwang, Ronan Le Bras, Antoine Bosselut, and Yejin Choi. 2020. [Back to the future: Unsupervised backprop-based decoding for counterfactual and abductive commonsense reasoning](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 794–805, Online. Association for Computational Linguistics.
- Lianhui Qin, Sean Welleck, Daniel Khashabi, and Yejin Choi. 2022. [COLD decoding: Energy-based constrained text generation with langevin dynamics](#). In *Advances in Neural Information Processing Systems*.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2022. Exploring the limits of transfer learning with a unified text-to-text transformer. *J. Mach. Learn. Res.*, 21(1).
- Andy Rosenbaum, Saleh Soltan, Wael Hamza, Yannick Versley, and Markus Boese. 2022. [LINGUIST: Language model instruction tuning to generate annotated utterances for intent classification and slot tagging](#). In *Proceedings of the 29th International Conference on Computational Linguistics*, pages 218–241, Gyeongju, Republic of Korea. International Committee on Computational Linguistics.

- Tianxiao Shen, Victor Quach, Regina Barzilay, and Tommi Jaakkola. 2020. [Blank language models](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 5186–5198, Online. Association for Computational Linguistics.
- Nisan Stiennon, Long Ouyang, Jeff Wu, Daniel M. Ziegler, Ryan Lowe, Chelsea Voss, Alec Radford, Dario Amodei, and Paul Christiano. 2020a. Learning to summarize from human feedback. In *Proceedings of the 34th International Conference on Neural Information Processing Systems, NIPS’20*, Red Hook, NY, USA. Curran Associates Inc.
- Nisan Stiennon, Long Ouyang, Jeffrey Wu, Daniel Ziegler, Ryan Lowe, Chelsea Voss, Alec Radford, Dario Amodei, and Paul F Christiano. 2020b. [Learning to summarize with human feedback](#). In *Advances in Neural Information Processing Systems*, volume 33, pages 3008–3021. Curran Associates, Inc.
- Yixuan Su, David Vandyke, Sihui Wang, Yimai Fang, and Nigel Collier. 2021. [Plan-then-generate: Controlled data-to-text generation via planning](#). In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 895–909, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Zewei Sun, Qingnan Jiang, Shujian Huang, Jun Cao, Shanbo Cheng, and Mingxuan Wang. 2022. [Zero-shot domain adaptation for neural machine translation with retrieved phrase-level prompts](#). *CoRR*, abs/2209.11409.
- Raymond Hendy Susanto, Shamil Chollampatt, and Lil-ing Tan. 2020. [Lexically constrained neural machine translation with Levenshtein transformer](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3536–3543, Online. Association for Computational Linguistics.
- Hongyin Tang, Miao Li, and Beihong Jin. 2019. [A topic augmented text generation model: Joint learning of semantics and structural features](#). 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 5090–5099, Hong Kong, China. Association for Computational Linguistics.
- Yi Tay, Mostafa Dehghani, Vinh Q. Tran, Xavier Garcia, Dara Bahri, Tal Schuster, Huaixiu Steven Zheng, Neil Houlsby, and Donald Metzler. 2022. [Unifying language learning paradigms](#). *CoRR*, abs/2205.05131.
- Ramakrishna Vedantam, C. Lawrence Zitnick, and Devi Parikh. 2015. Cider: Consensus-based image description evaluation. In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition (CVPR)*.
- Yufei Wang, Ian Wood, Stephen Wan, Mark Dras, and Mark Johnson. 2021. [Mention flags \(MF\): Constraining transformer-based text generators](#). 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 103–113, Online. Association for Computational Linguistics.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Remi Louf, Morgan Funtowicz, Joe Davison, Sam Shleifer, Patrick von Platen, Clara Ma, Yacine Jernite, Julien Plu, Canwen Xu, Teven Le Scao, Sylvain Gugger, Mariama Drame, Quentin Lhoest, and Alexander Rush. 2020. [Transformers: State-of-the-art natural language processing](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing: System Demonstrations*, pages 38–45, Online. Association for Computational Linguistics.
- Weijia Xu and Marine Carpuat. 2021. [EDITOR: An Edit-Based Transformer with Repositioning for Neural Machine Translation with Soft Lexical Constraints](#). *Transactions of the Association for Computational Linguistics*, 9:311–328.
- Kevin Yang and Dan Klein. 2021. [FUDGE: Controlled text generation with future discriminators](#). In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 3511–3535, Online. Association for Computational Linguistics.
- Ziyi Yang, Chenguang Zhu, Robert Gmyr, Michael Zeng, Xuedong Huang, and Eric Darve. 2020. [TED: A pretrained unsupervised summarization model with theme modeling and denoising](#). In *Findings of the Association for Computational Linguistics: EMNLP 2020*, pages 1865–1874, Online. Association for Computational Linguistics.
- Hanqing Zhang, Haolin Song, Shaoyu Li, Ming Zhou, and Dawei Song. 2022. [A survey of controllable text generation using transformer-based pre-trained language models](#). *CoRR*, abs/2201.05337.
- Tianyi Zhang*, Varsha Kishore*, Felix Wu*, Kilian Q. Weinberger, and Yoav Artzi. 2020. [Bertscore: Evaluating text generation with bert](#). In *International Conference on Learning Representations*.
- Yizhe Zhang, Guoyin Wang, Chunyuan Li, Zhe Gan, Chris Brockett, and Bill Dolan. 2020. [POINTER: Constrained progressive text generation via insertion-based generative pre-training](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 8649–8670, Online. Association for Computational Linguistics.
- Yusen Zhang, Yang Liu, Ziyi Yang, Yuwei Fang, Yulong Chen, Dragomir Radev, Chenguang Zhu, Michael

Zeng, and Rui Zhang. 2023. [MACSum: Controllable Summarization with Mixed Attributes](#). *Transactions of the Association for Computational Linguistics*, 11:787–803.

Wangchunshu Zhou, Yuchen Eleanor Jiang, Ethan Wilcox, Ryan Cotterell, and Mrinmaya Sachan. 2023. [Controlled text generation with natural language instructions](#). *CoRR*, abs/2304.14293.

Daniel M. Ziegler, Nisan Stiennon, Jeffrey Wu, Tom B. Brown, Alec Radford, Dario Amodei, Paul F. Christiano, and Geoffrey Irving. 2019. [Fine-tuning language models from human preferences](#). *CoRR*, abs/1909.08593.

A Recursive Algorithm

Algorithm 1 Recursive decoding

```

1: function RECURSIVEDECODING(exp)
2:   if not ContainNonterminal(exp) then
3:     return input
4:   end if
5:   if not ContainOptions(exp) then
6:     return Generate(exp)
7:   end if
8:   exp_before_opts, opts, exp_after_opts ←
   SplitByFirstOption(exp)
9:   for i, ch in enumerate(opts) do
10:    exp_ch ← exp_before_opts+ch
11:    opts[i] ← Generate(exp_ch)
12:   end for
13:   best_ch ← Generate(opts)
14:   remain_res ← Generate(exp_after_opts)
   return best_ch + remain_res
15: end function

```

B Dataset Statistics

Dataset	Train	Validation	Test
α NLG	50481	1780	3561
α NLI	169654	1532	3059
CommonGen	67216	993	-
Gigaword	-	189644	1933
IATE	-	-	414
Wiktionary	-	-	727
StoryCompletion	-	1871	1871

Table 9: Statistics of used dataset

C Generation Statistics

The data of average try and the first-time success rate during generation is presented in Table 10.

REI models tend to succeed on the first attempt for simple constraints, and only for challenging constraints, REI models would retry. Also, fine-tuned FLAN needs the least retry, while natural language instruction requires the most retry and may not be likely to succeed on the first try.

D Natural Language Instruction Examples

For the method of Natural Language Instruction on GPT-3.5, the instructions used on each task are shown in Table 11.

Task	Method					
	REI+FLAN-T5-xl		REI+GPT-3.5, 8 shot		NLI+GPT-3.5, 8 shot	
	Avg. Try	First SR.	Avg. Try	First SR.	Avg. Try	First SR.
aNLG, length	1.00	99.9	2.45	46.5	4.00	17.8
aNLG, lexicon	1.68	63.7	1.08	98.9	1.11	98.0
aNLG, length & lexicon	1.39	76.6	2.57	52.5	4.29	10.0
CommonGen	1.01	98.7	1.39	86.1	1.95	67.1
CommonGen, length	1.04	96.9	2.23	60.4	4.07	17.9
Gigaword, length	1.05	94.9	3.48	35.6	4.07	14.4
IATE	-	-	1.01	99.0	1.19	91.9
Wiktionary	-	-	1.05	97.0	1.08	92.8
StoryCloze, position	1.00	100.0	1.04	97.0	1.01	78.9

Table 10: Statistics of generation, presenting the average try (Avg. Try) and the first-time success rate (First SR.).

Task	Instruction Example
aNLG	The first sentence is " The Smiths were having holidays done of the children. " and the last sentence is " Ty's face lit up as he ran to the new toy, happily posing for photos. " . Insert a middle sentence with similar style, and the length shall not exceed 10 words.
aNLG, length	The first sentence is " The Smiths were having holidays done of the children. " and the last sentence is " Ty's face lit up as he ran to the new toy, happily posing for photos. " . Insert a middle sentence with similar style, and the length shall be exactly 7 words without counting punctuation.
aNLG, lexicon	The first sentence is " The Smiths were having holidays done of the children. " and the last sentence is " Ty's face lit up as he ran to the new toy, happily posing for photos. " . Insert a middle sentence with similar style, while also containing the keyword "bought".
aNLG, length & lexicon	The first sentence is " The Smiths were having holidays done of the children. " and the last sentence is " Ty's face lit up as he ran to the new toy, happily posing for photos. " . Insert a middle sentence with similar style, while also containing the keyword "bought", and the length shall be exactly 7 words without counting punctuation.
CommonGen	Generate a sentence that mentions all of these concepts in sequential order: "stood", "field", "looking".
CommonGen, length	Generate a sentence that mentions all of these concepts in sequential order with the word count of 10, punctuation ignored: "stood", "field", "looking".
Gigaword, length	Given the text "japan 's nec corp. and UNK computer corp. of the united states said wednesday they had agreed to join forces in supercomputer sales .", summarize the aforementioned text in a single phrase with the word count of 6.
IATE / Wiktionary	Translate from English to German using terminology "Interview":\n\nEnglish: That is what the Hollywood star has made abundantly clear in an interview.\nGerman:
StoryCloze, position	Given the first three sentences of the story "My friends all love to go to the club to dance. They think it's a lot of fun and always invite. I finally decided to tag along last Saturday." and two endings "My friends decided to keep inviting me out as I am so much fun." and "The next weekend, I was asked to please stay home.", infill the missing fourth sentence and choose the correct ending from the two.

Table 11: Examples of Natural Language Instructions