

Translate-and-Revise: Boosting Large Language Models for Constrained Translation

Pengcheng Huang^{1*}, Yongyu Mu^{1*}, Yuzhang Wu¹, Bei Li¹,
Chunyang Xiao³, Tong Xiao^{1,2†}, and Jingbo Zhu^{1,2}

¹NLP Lab, School of Computer Science and Engineering,
Northeastern University, Shenyang, China

²NiuTrans Research, Shenyang, China

³JP Morgan, United Kingdom

hpc1449181552@outlook.com lixiaoyumu9@gmail.com

{xiaotong, zhujingbo}@mail.neu.edu.cn

Abstract

Imposing constraints on machine translation systems presents a challenging issue because these systems are not trained to make use of constraints in generating adequate, fluent translations. In this paper, we leverage the capabilities of large language models (LLMs) for constrained translation, given that LLMs can easily adapt to this task by taking translation instructions and constraints as prompts. However, LLMs cannot always guarantee the adequacy of translation, and, in some cases, ignore the given constraints. This is in part because LLMs might be overly confident in their predictions, overriding the influence of the constraints. To overcome this overriding behaviour, we propose to add a revision process that encourages LLMs to correct the outputs by prompting them about the constraints that have not yet been met. We evaluate our approach on four constrained translation tasks, encompassing both lexical and structural constraints in multiple constraint domains. Experiments show 15% improvement in constraint-based translation accuracy over standard LLMs and the approach also significantly outperforms neural machine translation (NMT) state-of-the-art methods.

1 Introduction

Constrained translation seeks to generate translations that adhere to pre-specified constraints. To achieve this, conventional approaches impose constraints on machine translation systems and force them to follow the constraints during inference (Hokamp and Liu, 2017; Hasler et al., 2018; Dinu et al., 2019; Bergmanis and Pinnis, 2021b; Wang et al., 2022b; Ailem et al., 2022). More recently, large language models (LLMs) have been shown to be strong translation systems (Hendy et al., 2023; Moslem et al., 2023). They provide a general way to involve various instructions, demonstrations, and constraints into the translation process (Mu et al., 2023; Bogoychev and Chen, 2023), enabling us to perform constrained translation using off-the-shelf, well-trained LLMs.

While applying LLMs to constrained translation is straightforward, we observe empirically that even strong LLMs (i.e. GPT-3.5) do not always follow the instructions to obey constraints: LLMs’ predictions often override the guide of constraints, which result in missing constraints during translation. See Figure 1 for an example where we use an LLM to translate an English sentence to a Chinese sentence with a lexical constraint “COVID-19→新型冠状病毒”. We note that, despite significant effort in developing clear and instructive prompts, we were not able to improve the LLM in a single run of the LLM through the use of these constraints. For instance, we observed that when using open-source LLM to translate COVID-19, it tends to translate it as “新冠” more than 80% of the time, overlooking the constraint in the prompt to translate COVID-19 as “新型冠状病毒”. The problem consists of a real use case for what describes as ‘memo trap’ in the LLM literature (McKenzie et al., 2023).

To alleviate this problem and thus improve the accuracy to meet constraints, we propose to construct prompts iteratively that enable better focus on the unsatisfied constraints. The idea behind our approach

*Equal contribution.

†Corresponding author.

is to leverage the auto correction skills of LLMs by explicitly prompting them with which constraints are not satisfied (Madaan et al., 2023; Zhang et al., 2023c; Jiang et al., 2023). To do this, we introduce a revision step after the initial run of LLMs where we provide the LLMs with both the already-generated translation and the constraints that have not been covered. Then, we instruct the LLM to revise its output by taking these constraints into account.

We conduct experiments across four diverse constrained translation datasets, encompassing two distinct constraint types: lexical and structural. Our proposed “Translate-and-Revise” (TAR) approach consistently elevates the performance of LLMs in constrained translation, achieving state-of-the-art (SoTA) results on multiple datasets.

The contributions of this work are as follows:

- We introduce a novel TAR strategy that initially employs LLMs as constraint-aware translators and subsequently repropose them as revisers to revise translations that do not meet given constraints. We show that TAR significantly reduces missing constraints during translations.
- We rigorously evaluate our approach on four constrained translation datasets spanning multiple domains like news and electronics. Our results demonstrate a significant improvement in constraint fidelity and translation quality, outperforming existing methods and achieving SoTA results.
- To the best of our knowledge, our study is the first to evaluate LLMs across four distinct constrained translation datasets, thereby providing a robust LLM baseline for future research in the area. We believe our findings serve a solid baseline towards establishing more comprehensive benchmarks in the field of constrained translation.

2 Methods

Given a source language input and bilingual constraints, TAR first employs LLMs as translators for an initial translation. While this step often yields high-quality outputs, when the LLMs’ confidence during generation exceeds the guidance of the constraints, it results in suboptimal translation outputs. To mitigate the occurrence of missing constraints in LLMs-based translation, we introduce a reviser to enhance adherence to the constraints in the translation. The revision process is iterated multiple times until all constraints are satisfied, or the maximum allowable number of modifications is reached. The process of TAR is provided in Figure 1. Next, we describe TAR in more details.

2.1 Translate

Let $X = \{x_1, x_2, \dots, x_n\}$ be the source-language sentence with length n , and $Y = \{y_1, y_2, \dots, y_m\}$ be the target-language sentence with length m . The translation procedure can be written as:

$$Y = \text{Trans}(f(X)) \quad (1)$$

where $\text{Trans}(\cdot)$ symbolizes the translation model (either an NMT model or an LLM), and $f(\cdot)$ denotes a template by which we process X to make it suitable as the input of $\text{Trans}(\cdot)$.

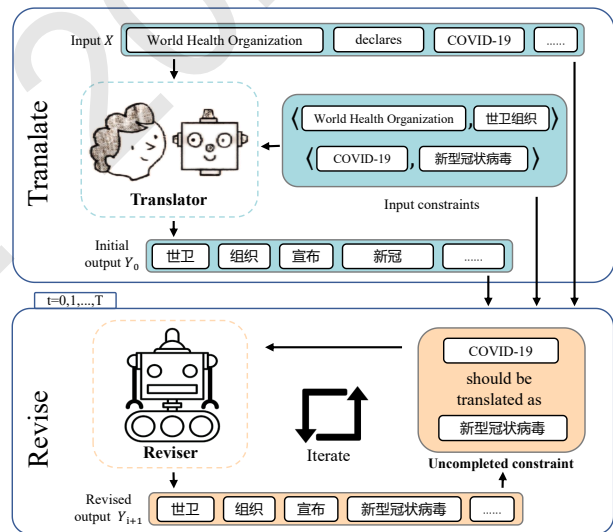


Figure 1: Given source language input X and constraint pairs, a *Translator* produces an initial translation Y_0 where COVID-19 is translated as “新冠”. Subsequently, a *Reviser* iteratively revise the translation Y_i to a better one Y_{i+1} , correctly translating COVID-19 as “新型冠状病毒”.

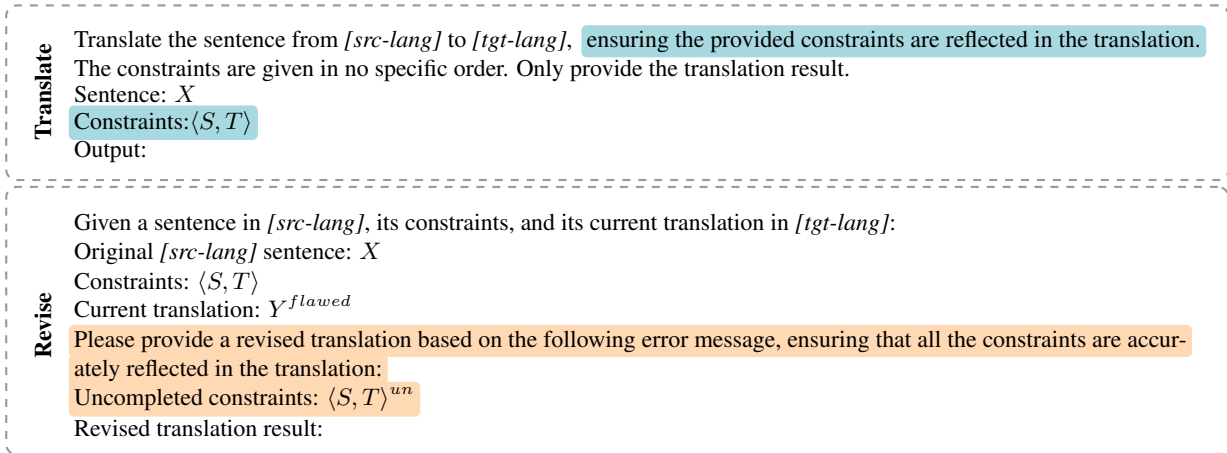


Figure 2: Two stages of TAR, in the Translate stage, constraints $\langle S, T \rangle$ are incorporated into the prompt to enable the model to generate preliminary translation results that meet the constraints to a certain extent. In the Revise stage, LLMs revise the flawed translation results Y^{flawed} with uncompleted constraints $\langle S, T \rangle^{un}$. The sections shaded in blue and yellow respectively represent the important parts of the two stages.

Let the $\langle S, T \rangle = \{\langle s_1, t_1 \rangle, \langle s_2, t_2 \rangle, \dots, \langle s_k, t_k \rangle\}$ represents the bilingual constraints with k pairs in total. Constrained translation needs the system to accurately translate each source constraint s_i to its corresponding target constraint t_i . This process can be represented by the following equation:

$$Y = \text{Trans}(f(X, \langle S, T \rangle)) \quad (2)$$

$s.t. T \in Y$

Since conventional translation instructions never impose constraints on LLMs, they frequently fall short of satisfying constraints. In this work, we propose to integrate these constraints directly into the prompts and employ an instruction based on natural language specifically tailored for constrained translation tasks. Our template $f(\cdot)$ is shown in Figure 2 which can effectively turn LLMs into constraint-aware translators.

2.2 Revise

However, the LLM-based translation cannot always cover all original constraints. We randomly sampled 20 incorrect translation results and observed that, in datasets like WMT21 Terminology Translation (Alam et al., 2021), to 95%(19/20) of the cases, the tokens generated by the model were similar to the expected constraints meaning and exhibited high confidence levels. The confidence level of LLMs in generating these tokens remained virtually unchanged, whether or not constraints were included in the instructions, revealing overconfidence in generation while overlooking the constraints.

We notice a strong connection between our real use case and ‘memo trap’ (McKenzie et al., 2023) as unsatisfied constraints often pertain to non-mainstream translations resulting terms used with lower frequency and the incorrect translations usually refer to the mainstream translations. Compared ‘memo trap’, we show that the phenomenon extends to toy settings and is prominent even for real applications and for SOTA models like GPT-3.5.

To overcome these challenges, we initially employ a rule-based method to identify which constraints are not completed. Subsequently, these uncompleted constraints $\langle S, T \rangle^{un}$, along with the source language input X , flawed translations output Y^{flawed} , and all other given constraints $\langle S, T \rangle$, are passed to the LLM. At this juncture, the LLM assumes the role of a reviewer, tasked with revising flawed translation upon receipt of uncompleted constraints. The aforementioned process is defined by the following formula:

$$Y = \text{Revise}(f(X, \langle S, T \rangle, \langle S, T \rangle^{un}, Y^{flawed})) \quad (3)$$

where $\text{Revise}(\cdot)$ symbolizes the reviser. Furthermore, TAR can continuously iterate in the loop of detecting uncompleted constraints and making revisions until a stopping condition is met. This condition is either the iteration i reaches a specified count or the translation satisfies all constraints. We represent this iterative process as follows:

$$Y_{i+1} = \text{Revise}(f(X, \langle S, T \rangle, \langle S, T \rangle_i^{un}, Y_i)) \quad (4)$$

where the translation Y_i from the previous iteration, combined with its uncompleted constraints $\langle S, T \rangle_i^{un}$ and other inputs, is sent into the reviser to produce a more precise translation Y_{i+1} . By highlighting uncompleted translation constraints and comparing flawed translation results, human translators are able to satisfy these constraints and optimize translation output. Empirically, we find that LLMs can revise the translation results similarly to human translators, while being more efficient and cost-effective. Additionally, we discuss the time and financial costs of multiple iterations in Appendix A.

3 Experiments

In this study, we evaluate the performance of the TAR in constrained translation. While most previous research has typically focused on just one or two constrained translation tasks (Dinu et al., 2019; Wang et al., 2022a; Hashimoto et al., 2019; Zeng et al., 2023), our evaluation expands to two types of constraints: lexical constraints and structural constraints, covering four practical scenarios: general lexically constrained translation, translation with terminology constraints, translation with named entity constraints, and structured document translation.

3.1 Setup

Datasets Detailed information about the datasets we used can be found in Table 1. Lexical constraints refer to sentences with predefined word or phrase constraints sourced from existing databases. Structural constraints, in contrast, encompass inline markup tag constraints like XML tags, for example, $\langle \text{ph} \rangle$ and $\langle / \text{ph} \rangle$. Here are more details about the datasets that we use in this work:

General Lexically Constrained Translation: This is based on a dataset⁰ provided by (Dinu et al., 2019). The dataset is derived from *newstest2017* En \rightarrow De. Lexical constraints are extracted with guidance from two general-domain term databases: IATE and Wiktionary¹.

Terminology Translation: To benchmark against SoTA NMT systems, we employ the official test set from the *terminology translation* task in WMT21²(WMT21 TT) (Alam et al., 2021).

Entity Translation: We also endeavor to evaluate our method using the extensive Entity Translation Corpus (ETC) (Zeng et al., 2023), which comprises six test sets from the WMT News Translation Task spanning 2015-2021. For alignment, we employ spaCy NER models³ to extract source entities and use awesome-align (Dou and Neubig, 2021) for their correspondence.

Structured Document Translation: Following recent works (Wang et al., 2022a), we conduct experiments on the LXM dataset⁴ (Hashimoto et al., 2019), in which XML tags are hierarchically distributed

Corpus	Language direction	#Sent.	#Lines with const.	#Const.
<i>Lexical constraint</i>				
IATE Wiktionary	En-De	2000	414	452
	En-De	2000	727	884
WMT21 TT	En-Ru	2100	1307	2524
	En-Zh	2100	1191	2229
ETC	En-Zh	19144	12040	35253
	En-Ru	12985	3917	10308
	Zh-En	19144	12040	35253
	Ru-En	12985	3917	10308
<i>Structural constraint</i>				
LXM	En-Zh	2000	518	884
	En-De	2000	520	942
	En-Ru	2000	554	993
	En-Fr	2000	575	1051

Table 1: Statistics of the datasets we used for four different task: the total number of sentences in datasets ($\#Sent.$), the number of lines with constraints ($\#Lines\ with\ const.$), the number of constraints ($\#const.$).

⁰https://github.com/mtresearcher/terminology_dataset

¹Available at <https://iate.europa.eu/home> and <https://www.wiktionary.org/>.

²<https://www.statmt.org/wmt21/terminology-task.html>

³<https://pypi.org/project/spacy/>

⁴<https://github.com/salesforce/localization-xml-mt>

Method	BLEU	CCR%	BLEU	CCR%
Dataset	IATE		Wiktionary	
Transformer	25.8	76.3	26.0	76.9
Const. Dec.	25.3	82.0	25.8	99.5
Code-switching	26.0	94.5	26.3	93.4
Append	26.0	92.8	26.9	90.7
RTT	27.2	99.6	27.8	98.3
LLM Trans.	32.0	85.4	32.0	88.7
LLM Const. Trans.	32.0	96.2	32.1	97.2
+Revision	32.0 _(+0.0)	98.9 _(+2.7)	32.0 _(-0.1)	98.9 _(+1.7)

Table 2: Results of the general lexically constrained translation task. The **highest** scores among the various systems are highlighted in bold, while the *second-best* scores are emphasized in italics for clarity.

throughout the source and target text.

Evaluation Metrics Consistent with previous studies (Dinu et al., 2019; Ailem et al., 2022; Wang et al., 2022a; Zeng et al., 2023), we employ BLEU (Papineni et al., 2002) and the constraints completion rate (CCR) to assess translation quality and constraint-based translation accuracy except for structured document translation. For entity translation, we also incorporate the COMET score⁵ (Rei et al., 2020) for comparison with the work of Zeng (2023). In structured document translation, we utilize sacreBLEU (Post, 2018) to compute the XML-based BLEU score⁶. Additionally, we measure structured constraint-based translation accuracy using the structure accuracy rate (SAR) and structure match rate (SMR). Here, SAR evaluates the compatibility of translation results with XML parsers, while SMR ensures the translated XML structure aligns with the reference. Both metrics are assessed using lxml⁷.

Baselines To ensure thorough evaluation, apart from comparing TAR with the LLM baseline without revision, we also compare TAR with representative methods across each of the four tasks. For general lexically constrained translation, our baselines include the vanilla Transformer (Vaswani et al., 2017), Const.Dec. (Post and Vilar, 2018), Code-switching (Dinu et al., 2019), Append (Dinu et al., 2019), and Robust Terminology Translation (RTT) (Zhang et al., 2023b). For the terminology translation task, our baselines are derived from the top three submissions of WMT21. They include HW-TSC (Wang et al., 2021b), Term-Mind-sys2 (Wang et al., 2021a), ProMT.soft (Molchanov et al., 2021), TildeMT (Bergmanis and Pinnis, 2021a), and Lingua Custodia (Ailem et al., 2021). For entity translation, we utilize the vanilla Transformer, Code-switching, Placeholder (Yan et al., 2019), and Extract and Attend (Zeng et al., 2023). Structured document translation baselines consist of the vanilla Transformer, Split-Inject (Al-Anzi et al., 1997), and Template (Wang et al., 2022a) methods.

Model Configurations We initially investigate the potential of LLMs to act as both the translator and reviser within TAR. Our primary choice for LLMs is gpt-3.5-turbo-0613⁸, chosen for its exceptional translation capabilities and proficiency in adhering to instructions. Additionally, we assess its revising process for NMT models in Section 4.1 and explore whether TAR provides consistent improvements across different LLMs in Section 4.2. The decoding parameters for these models remain at their default settings, except for the sampling temperature, which is set to 0. We employ natural language-based prompts in a one-shot manner, merging uncompleted constraints with the source language, flawed translation results, and original constraints to form the reviser’s input. These prompts are depicted in Figure 2.

Detection of Uncompleted Constraints To identify unmet constraints in translations, we employ a rule-based procedure that leverages scripts designed for calculating CCR. This procedure assesses how well the translation adheres to the constraints. We further explore the capacity of LLMs to autonomously verify constraint completion and offer detailed feedback to the reviser in Section 3.3.

⁵wmt22-comet-da

⁶XML tags are treated as an integral part of the sentences during BLEU score calculation.

⁷<https://lxml.de/>

⁸<https://openai.com/>

Method	BLEU	CCR%	BLEU	CCR%
Direction	English-Chinese		English-Russian	
HW-TSC	40.7	88.6	-	-
TermMind-sys2	40.5	85.6	-	-
ProMT.soft	-	-	31.1	90.9
TildeMT	-	-	28.2	86.3
Lingua Custodia	29.6	82.8	28.8	85.4
LLM Trans.	36.3	87.2	29.7	85.9
LLM Const. Trans.	36.4	92.6	30.1	95.8
+Revision	35.9 _(-0.5)	95.9 _(+3.3)	30.3 _(+0.2)	97.5 _(+1.7)

Table 3: Results of the terminology translation task for both English-Chinese and English-Russian.

Method	BLEU	COMET	CCR%	BLEU	COMET	CCR%
<i>Direction</i>	<i>English-Chinese</i>			<i>Chinese-English</i>		
Transformer	26.3	34.8	57.3	27.5	41.5	59.0
Code-switching	25.9	41.4	70.5	27.2	45.0	71.1
Placeholder	26.4	42.9	71.4	27.5	47.2	72.1
Extract & Attend	26.8	48.6	72.3	28.0	50.1	72.5
LLM Trans.	39.5	87.6	77.2	27.3	83.8	85.3
LLM Const. Trans.	40.0	87.4	96.6	28.9	83.5	94.2
+Revision	40.0 _(+0.0)	87.4 _(+0.0)	97.6 _(+1.0)	28.9 _(+0.0)	83.5 _(+0.0)	97.3 _(+3.1)
<i>Direction</i>	<i>English-Russian</i>			<i>Russian-English</i>		
Transformer	31.8	52.2	40.0	34.6	54.0	48.7
Code-switching	30.5	55.2	50.4	32.0	56.7	50.2
Placeholder	31.9	57.6	50.3	34.7	59.1	50.7
Extract & Attend	32.7	62.2	57.3	35.4	63.5	58.4
LLM Trans.	31.8	89.9	64.5	36.0	85.8	76.8
LLM Const. Trans.	32.6	89.8	88.8	36.8	85.8	96.7
+Revision	32.5 _(-0.1)	89.8 _(+0.0)	89.8 _(+1.0)	36.8 _(+0.0)	85.8 _(+0.0)	97.5 _(+0.8)

Table 4: Results of the entity translation task for English-Chinese, English-Russian, Chinese-English and Russian-English.

Method	BLEU	SAR%	SMR%	BLEU	SAR%	SMR%
<i>Direction</i>	<i>English-Chinese</i>			<i>English-German</i>		
Transformer	61.2	99.85	99.25	52.7	99.80	99.20
Split-Inject	57.0	100.00	99.30	50.7	100.00	99.80
Template	61.5	100.00	99.80	53.6	100.00	99.80
LLM Trans.	55.1	99.95	98.95	49.2	99.95	99.25
LLM Const. Trans.	56.4	100.00	99.50	49.2	99.95	99.25
+Revision	56.5 _(+0.1)	100.00 _(+0.00)	99.75 _(+0.25)	49.2 _(+0.0)	100.00 _(+0.05)	99.30 _(+0.05)
<i>Direction</i>	<i>English-French</i>			<i>English-Russian</i>		
Transformer	65.3	99.55	99.30	44.9	99.45	98.90
Split-Inject	66.1	100.00	100.00	43.1	100.00	99.85
Template	67.3	100.00	100.00	45.8	100.00	99.80
LLM Trans.	58.1	99.90	99.30	34.4	99.90	99.35
LLM Const. Trans.	59.3	100.00	99.95	36.0	100.00	99.60
+Revision	59.3 _(+0.0)	100.00 _(+0.00)	99.95 _(+0.00)	36.0 _(+0.0)	100.00 _(+0.00)	99.75 _(+0.15)

Table 5: Results of the structured document translation for English-Chinese, English-German, English-Chinese and English-Russian.

3.2 Main results

Table 2, Table 3, Table 4, and Table 5 detail the performance of TAR on general lexically constrained translation, terminology translation, entity translation, and structured document translation, respectively. Here are our main results.

Comparison with base LLMs TAR consistently boosts the performance of LLMs in constrained translation. Two primary factors contribute to this improvement:

(1) Our natural language-based prompts, as opposed to the conventional few-shot translation prompts (Hendy et al., 2023), are more effective for constrained translation. Specifically, in terminology translation (refer to Table 3), our prompts lead to an average BLEU score increase of 0.3 and a CCR rise of 7.7%.

We noticed significant gains over base LLMs in entity translation. Across four language directions, there is a consistent uplift in BLEU scores, averaging an increase of 0.9. Notably, the CCR experiences increases of 19.4%, 8.9%, 24.3%, and 19.9% respectively. This marked improvement can primarily be attributed to the superior instruction-following capabilities of LLMs. By incorporating constraints in instructions, we can guide the model more effectively to address these constraints, alleviating the issue of LLMs struggling to correctly translate named entities.

(2) Revision effectively improves constraint-based translation accuracy across all datasets without sacrificing translation quality, addressing the tendency of LLM-based translations to overlook constraints.

As illustrated in Table 6, we observe that the LLM exhibits overconfidence in its prediction overriding the influence of constraints, especially when the constraint suggests an uncommon translation for a polysemous word. For instance, the LLM translates “COVID-19” to the more commonly used “新冠” instead of adhering to the target constraint “新型冠状病毒”. We speculate this overconfidence in LLMs stems from their greater exposure to “新冠” compared to “新型冠状病毒” during the pre-training phase, which may lead them to be overly loyal to certain patterns, thereby preventing them from meeting certain constraints.

However, the revision step can effectively cut down the possibility of missing constraints by explicitly prompting LLMs with which constraints are not satisfied. Experimental results indicate that our revision strategy led to average improvements in CCR by 2.2% for lexically constrained translation, 2.5% for terminology translation, and 1.5% for entity translation across various language directions. While the improvements in SCR and SMR for structured document translation might not seem prominent, it’s primarily because the initial translation is already at a 100% performance.

Comparison with supervised methods The SoTA methods for these four constrained translation datasets predominantly rely on pseudo-data augmentation. Through our experimental results, we observe that these methods nearly achieve perfection on the IATE, Wiktionary, and LXM datasets (as evidenced in Table 2 and Table 5). We contend that the test sets for these datasets might be relatively straightforward, and the constraints they encompass are frequently encountered in training sets. Therefore, they may not accurately reflect real-world applications where constraints might span multiple domains and are infrequently seen in the training data. However, when assessed on the ETC (as depicted in Table 4), which comprises test data spanning from 2015 to 2021, showcasing a rich diversity in constraint domains, the efficacy of traditional data augmentation methods seems to be poor. In contrast, TAR’s performance remains stable, demonstrating comparable constraint-based translation accuracy on ETC as with other datasets. This highlights TAR’s proficiency in handling the diverse constraint requirements found in real-world situations.

3.3 Impact of Inputs on Reviser Performance

The reviser receives inputs including the source language sentence, translation results, given constraints, and the uncompleted constraints. To evaluate the significance of each component, we conducted experiments wherein we omitted specific elements from the input. The variations include: 1) Excluding uncompleted constraints; 2) Excluding original constraints; and 3) Only indicating to the model that the translation is flawed without specifying the uncompleted constraints. All other settings remain unchanged. The comparative outcomes on IATE and Wiktionary are presented in Table 7.

From our observations, the CCR scores of the variants show a decline compared to the default input of the TAR reviser on both datasets. Interestingly, the omission of the original constraints has a more pronounced impact on CCR. This could be attributed to the negative modification of completed constraints made by the reviser when it is inaccessible to the given constraints. Additionally, when both elements

Constraints	(WHO, 世界卫生组织); (COVID-19, 新型冠状病毒)
Source	On 11 March 2020, WHO characterized COVID-19 as a pandemic.
Reference	2020年3月11日 世界卫生组织 将 新型冠状病毒 列为大流行病。
Const. Trans.	2020年3月11日 世界卫生组织 将新冠确定为大流行病。
+ Revision	2020年3月11日 世界卫生组织 将 新型冠状病毒 定性为大流行病。

Table 6: A case study of TAR: Initially, the translator rendered “COVID-19” as the more prevalent “新冠” in Chinese. With the intervention of the reviser, it was accurately translated as “新型冠状病毒”, thereby satisfying all constraints.

Setting	IATE		Wiktionary	
	BLEU	CCR%	BLEU	CCR%
base	32.0	96.2	32.1	97.2
after revise	32.0	98.9	32.1	98.9
- Uncompleted const.	32.0	97.4	32.1	97.9
- Original const.	32.0	97.1	32.0	97.7
- Both	32.0	96.2	32.1	97.2
+ Detected by LLM	31.9	95.8	32.0	97.2

Table 7: BLEU and CCR scores of ablation on supplementary feedback. “Uncompleted constraints” and “Original constraints” are parts of the input received by the reviser.

are excluded, there’s no noticeable difference in the CCR before and after the revision process.

Furthermore, as shown by ‘+ Detected by LLM’ in Table 7, using LLMs to detect uncompleted constraints and then feeding them back to the reviser may degrade translation performance. Further analysis of the results reveals challenges faced by LLMs inaccurately identifying unsatisfied constraints; they often mistakenly believe certain constraints have been met. Such inaccurate feedback not only fails to enhance the quality of the translation but might even deteriorate it. These insights emphasize the criticality of supplying the reviser with exact and thorough constraint information.

4 Analysis

4.1 TAR Augments NMT Translators

In our study, we initially depended on constraint-aware translators to produce preliminary translation results. However, in real-world scenarios, industry practitioners often possess powerful domain-agnostic NMT models. These models, due to their lack of training with specific constraints, frequently fall short in constrained translation tasks. In this section, we integrate TAR into these general-purpose NMT models. By iteratively optimizing the NMT translation results through TAR, we can significantly enhance the CCR of the translation while ensuring its quality.

Specifically, we first the WMT21 champion model (Tran et al., 2021) to obtain a preliminary translation result. Since this model is not specifically trained for constraints, the initial translation often exhibits a suboptimal CCR. Building on this, we apply TAR to revise this outcome, iteratively optimizing to form the final translation result.

The experimental results on WMT21 TT and ETC datasets are presented in Table 8, we can see that TAR bolsters both BLEU and CCR scores of NMT models. On average, we observed an uplift of 0.7 and 0.8 in BLEU scores, coupled with impressive gains of 10.7% and 14.3% in CCR across the two datasets and language pairs. Although there remains a gap in constraint-based translation accuracy compared to the standard TAR, it generally exhibits superior translation quality. This insight demonstrates that when equipped with TAR, even domain-agnostic NMT models can adeptly tackle constrained translation. This eliminates the need for forced decoding algorithms or additional training, greatly enhancing their usability in constrained translation applications.

4.2 Scaling TAR to More LLMs

To evaluate the scalability and robustness of TAR across different models, we applied it to a variety of LLMs, including commercial models like GPT-3 and GPT-4, as well as the open-source Qwen⁹ (Bai et al., 2023). We maintained consistency in all other settings. Experiments were conducted in the “En-Zh” and “En-Ru” language directions for terminology translation. As shown in Figure 3, TAR consistently improves the performance of various LLMs in constrained translation tasks. To be specific, the average CCR for GPT-3, GPT-4, and Qwen increases by 20%, 11.7%, and 14.4%, respectively. Comparing the results of ChatGPT with those of GPT-3 and GPT-4, it’s evident that TAR enables more powerful models to fully harness their capabilities in constrained translation. Intriguingly, although the CCR score

	BLEU	CCR%	BLEU	CCR%
WMT21 TT	<i>English-Chinese</i>		<i>English-Russian</i>	
TAR	35.9	95.9	30.3	97.5
NMT	34.5	85.6	33.6	85.3
+ Revision	35.2 _(+0.7)	95.6 _(+10.0)	34.0 _(+0.4)	96.6 _(11.3)
ETC	<i>English-Chinese</i>		<i>English-Russian</i>	
TAR	40.0	97.6	32.5	89.8
NMT	41.2	74.5	40.1	72.3
+ Revision	42.7 _(+1.5)	92.8 _(+18.3)	40.2 _(+0.1)	82.4 _(+10.1)

Table 8: Results of applying TAR to NMT on the WMT21 TT and ETC datasets for both English-Chinese and English-Russian.

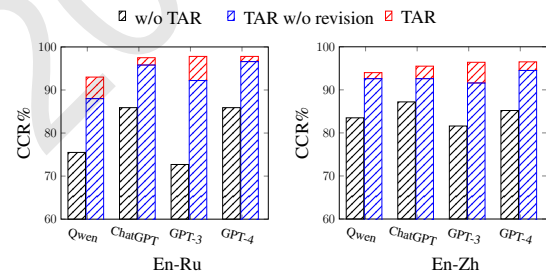


Figure 3: TAR results on WMT21 TT using Qwen, ChatGPT, GPT-3 (text-davinci-003) and GPT-4. Here “w/o TAR” represents the use of the conventional translation prompt. “TAR w/o revision” indicates the use of a prompt with constraints, but without reviser. Meanwhile, “TAR” denotes the full method that includes revisions.

⁹Qwen-14B-Chat

of GPT-3 in the initial translation substantially trails that of ChatGPT, it surpasses ChatGPT post-revision. While the performance of Qwen lags slightly behind ChatGPT, the improvement brought by TAR is still notable.

4.3 Revision Iterative Round and Prompt Ensemble

The reviser, in its function, takes the translation result and uncompleted constraints as input. Naturally, one might consider iteratively revising the output multiple times. The question arises: how many iterations strike the optimal balance? Here, we employed constraint-aware LLMs and NMT on the “En-Zh” and “En-Ru” directions of the terminology translation dataset. We assessed performance across different iterative rounds of the revision module, consistently using the same prompts for each iterative phase.

As presented in Figure 4a, there is a significant leap in performance predominantly during the initial revision. Although performance does enhance with increasing iterations, the rate of improvement starts to taper off, indicating diminishing returns.

Furthermore, we also investigated whether the LLM reviser can benefit from varying prompts in multiple iterations. To experiment this, after designing several revision templates, we randomly select one for each iterative round. Thus, a complete multi-round revision process can utilize various templates. This design is in part similar to prompt ensemble methods (Zhang et al., 2023a; Pitis et al., 2023), combining benefits of various prompts, akin to ensemble learning (Dong et al., 2020). Experimental results are shown in Figure 4b. Compared to applying a single template in revision iterations, utilizing diverse templates achieves superior performance.

4.4 More Analysis

Due to space limitations, we provide a more detailed analysis of our method in the appendix, including the additional costs incurred by TAR, the impact on performance as the number of constraints increases, reasons for potential declines in BLEU scores during the revision phase for certain datasets, and the performance of traditional constrained translation data augmentation methods on LLMs.

5 Related Work

5.1 Constrained Translation

Machine translation has made considerable progress in incorporating pre-specified constraint, which can be categorized into hard constrained translation and soft constrained translation.

Hard constrained translation: This line of research expanded the original search space via decoding algorithm modification to strictly incorporate constraints (Hokamp and Liu, 2017; Post and Vilar, 2018; Hu et al., 2019). However, while these methods achieve a high constraint-based translation accuracy, they tend to be computationally expensive and can sometimes compromise translation quality (Hasler et al., 2018; Zhang et al., 2021).

Soft constrained translation: Here, research primarily centers on data augmentation strategies to train NMT models to integrate constraints. Several techniques have been proposed, including: replacing source language constraints with special token (Crego et al., 2016; Wang et al., 2017; Zhang et al., 2023b); substituting source language constraints with target language constraints (Song et al., 2019; Dinu et al., 2019); using inline annotations to individually mark source and target language constraints (Ailem




Figure 4: (a) Improvements in CCR with each iteration. (b) Red denotes consistent template use across three iterations, while blue indicates alternating templates.

et al., 2022; Chen et al., 2021b); and (Wang et al., 2022a) employing a template to transform the constrained translation into constraint reordering. There are also studies that modified the model architecture to better integrate vectorized constraints representation (Li et al., 2020; Wang et al., 2022b), alignment information (Chen et al., 2021a). These approaches heavily rely on data quality or necessitate structural modifications, limiting their practicality. Moreover, they often falter when addressing diverse real-world requirements. In contrast, TAR doesn't require training on constraint-specific data and is adept at handling varied constraint scenarios.

5.2 Automatic Post-editing

Several studies have aimed to develop neural-based models for automatic post-editing (APE) in translation (Vu and Haffari, 2018; Correia and Martins, 2019; Shterionov et al., 2020). Chatterjee(2019) investigated the application of deep learning techniques for APE and introduced novel architectures to improve the quality of post-edited translations. Góis(2020) examined the application of automated ordering methods to improve translations. Voita(2019) introduced a context-aware approach to APE, integrating source context information into the neural framework to produce improved post-edits. Cholampatt(2020) investigated the application of APE in enhancing the translation performance of NMT models. These methods primarily focus on enhancing the overall translation quality. However, it's crucial to understand that not all words within a sentence carry equal importance. The precise translation of terminologies and entities significantly impacts user experience. Our proposed TAR specifically addresses the challenge of ensuring more accurate translations for these constraints.

6 Conclusion

In this work, we introduce the TAR prompting method, adeptly leverages LLMs for constrained translation. Our approach involves a two-step process: first using LLMs for constrained translation, and subsequently deploying them to revise translations with uncompleted constraints. Our approach mainly improves the constraint accuracy while maintaining translation quality by overcoming the 'memo trap' from the LLMs during translation using dedicated revision prompts in an iterative manner.

We further show that TAR can be applied to LLM based translation systems as well as traditional NMT systems, in both cases resulting in better constraint accuracy while maintaining translation quality and the technology is not limited to particular LLMs. More generally, our study sheds light on the importance of accurate feedback in general for LLM revision to work effectively.

7 Limitations

While we have demonstrated TAR's efficacy across four constrained translation datasets, real-world applications are considerably more varied, our prompts might not always yield optimal outcomes. In fact, the essence of TAR lies in its revision mechanism. However, as emphasized in Section 3.3, detecting constraint adherence using LLMs poses challenges. Rule-based methods, though effective in offering accurate feedback to the reviser, can falter in broader constraint scenarios, such as controlled text generation demanding specific stylistic alignment. In such contexts, devising a method to secure accurate and efficient feedback to guide model revisions remains a research imperative. We believe that overcoming these challenges will solidify TAR's standing as a universally effective framework across diverse constraint scenarios.

Acknowledgements

This work was supported in part by the National Science Foundation of China (No.62276056), the Natural Science Foundation of Liaoning Province of China (2022-KF-16-01), the Fundamental Research Funds for the Central Universities (Nos. N2216016 and N2316002), the Yunnan Fundamental Research Projects (No. 202401BC070021), and the Program of Introducing Talents of Discipline to Universities, Plan 111 (No.B16009).

References

- Melissa Ailem, Jingshu Liu, and Raheel Qader. 2021. Lingua custodia’s participation at the wmt 2021 machine translation using terminologies shared task. In *Proceedings of the Sixth Conference on Machine Translation*, pages 799–803.
- Melissa Ailem, Jingshu Liu, and Raheel Qader. 2022. Encouraging neural machine translation to satisfy terminology constraints. In *Actes de la 29e Conférence sur le Traitement Automatique des Langues Naturelles. Volume 1 : conférence principale*, page 446.
- F Al-Anzi, K Al-Zame, M Husain, and H Al-Mutairi. 1997. Automatic english/arabic html home page translation tool. In *Proc. 1st Workshop Technol. Arabizing Internet*.
- Md Mahfuz Ibn Alam, Ivana Kvapilíková, Antonios Anastasopoulos, Laurent Besacier, Georgiana Dinu, Marcello Federico, Matthias Gallé, Kweon Woo Jung, Philipp Koehn, and Vassilina Nikoulina. 2021. Findings of the WMT shared task on machine translation using terminologies. In Loïc Barrault, Ondrej Bojar, Fethi Bougares, Rajen Chatterjee, Marta R. Costa-jussà, Christian Federmann, Mark Fishel, Alexander Fraser, Markus Freitag, Yvette Graham, Roman Grundkiewicz, Paco Guzman, Barry Haddow, Matthias Huck, Antonio Jimeno-Yepes, Philipp Koehn, Tom Kocmi, André Martins, Makoto Morishita, and Christof Monz, editors, *Proceedings of the Sixth Conference on Machine Translation, WMT@EMNLP 2021, Online Event, November 10-11, 2021*, pages 652–663. Association for Computational Linguistics.
- Jinze Bai, Shuai Bai, Yunfei Chu, Zeyu Cui, Kai Dang, Xiaodong Deng, Yang Fan, Wenbin Ge, Yu Han, Fei Huang, Binyuan Hui, Luo Ji, Mei Li, Junyang Lin, Runji Lin, Dayiheng Liu, Gao Liu, Chengqiang Lu, Keming Lu, Jianxin Ma, Rui Men, Xingzhang Ren, Xuancheng Ren, Chuanqi Tan, Sinan Tan, Jianhong Tu, Peng Wang, Shijie Wang, Wei Wang, Shengguang Wu, Benfeng Xu, Jin Xu, An Yang, Hao Yang, Jian Yang, Shusheng Yang, Yang Yao, Bowen Yu, Hongyi Yuan, Zheng Yuan, Jianwei Zhang, Xingxuan Zhang, Yichang Zhang, Zhenru Zhang, Chang Zhou, Jingren Zhou, Xiaohuan Zhou, and Tianhang Zhu. 2023. Qwen technical report.
- Toms Bergmanis and Mārcis Pinnis. 2021a. Dynamic terminology integration for covid-19 and other emerging domains. In *Proceedings of the Sixth Conference on Machine Translation*, pages 821–827.
- Toms Bergmanis and Mārcis Pinnis. 2021b. Facilitating terminology translation with target lemma annotations. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 3105–3111.
- Nikolay Bogoychev and Pinzhen Chen. 2023. Terminology-aware translation with constrained decoding and large language model prompting. *arXiv preprint arXiv:2310.05824*.
- Rajen Chatterjee. 2019. Automatic post-editing for machine translation. *arXiv preprint arXiv:1910.08592*.
- Guanhua Chen, Yun Chen, and Victor OK Li. 2021a. Lexically constrained neural machine translation with explicit alignment guidance. In *Proceedings of the AAAI Conference on Artificial Intelligence*, volume 35, pages 12630–12638.
- Guanhua Chen, Yun Chen, Yong Wang, and Victor OK Li. 2021b. Lexical-constraint-aware neural machine translation via data augmentation. In *Proceedings of the Twenty-Ninth International Conference on International Joint Conferences on Artificial Intelligence*, pages 3587–3593.
- Shamil Chollampatt, Raymond Hendy Susanto, Liling Tan, and Ewa Szymanska. 2020. Can automatic post-editing improve nmt? In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*, pages 2736–2746.
- Gonçalo M Correia and André FT Martins. 2019. A simple and effective approach to automatic post-editing with transfer learning. In *Proceedings of the 57th Conference of the Association for Computational Linguistics*, pages 3050–3056.
- Josep Crego, Jungi Kim, Guillaume Klein, Anabel Rebollo, Kathy Yang, Jean Senellart, Egor Akhanov, Patrice Brunelle, Aurelien Coquard, Yongchao Deng, et al. 2016. Systran’s pure neural machine translation systems. *arXiv preprint arXiv:1610.05540*.
- Georgiana Dinu, Prashant Mathur, Marcello Federico, and Yaser Al-Onaizan. 2019. Training neural machine translation to apply terminology constraints. In *Proceedings of the 57th Conference of the Association for Computational Linguistics*, pages 3063–3068.
- Xibin Dong, Zhiwen Yu, Wenming Cao, Yifan Shi, and Qianli Ma. 2020. A survey on ensemble learning. *Frontiers Comput. Sci.*, 14(2):241–258.

- Zi-Yi Dou and Graham Neubig. 2021. Word alignment by fine-tuning embeddings on parallel corpora. In *Proceedings of the 16th Conference of the European Chapter of the Association for Computational Linguistics: Main Volume*, pages 2112–2128.
- António Góis, Kyunghyun Cho, and André Martins. 2020. Learning non-monotonic automatic post-editing of translations from human orderings. In *Proceedings of the 22nd Annual Conference of the European Association for Machine Translation*, pages 205–214.
- Kazuma Hashimoto, Raffaella Buschiazzo, James Bradbury, Teresa Marshall, Richard Socher, and Caiming Xiong. 2019. A high-quality multilingual dataset for structured documentation translation. In *Proceedings of the Fourth Conference on Machine Translation*, pages 116–127.
- Eva Hasler, Adrià De Gispert, Gonzalo Iglesias, and Bill Byrne. 2018. Neural machine translation decoding with terminology constraints. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 506–512.
- Amr Hendy, Mohamed Abdelrehim, Amr Sharaf, Vikas Raunak, Mohamed Gabr, Hitokazu Matsushita, Young Jin Kim, Mohamed Afify, and Hany Hassan Awadalla. 2023. How good are gpt models at machine translation? a comprehensive evaluation. *arXiv preprint arXiv:2302.09210*.
- Chris Hokamp and Qun Liu. 2017. Lexically constrained decoding for sequence generation using grid beam search. In *Proceedings of the 55th Annual Meeting of the Association for Computational Linguistics*, pages 1535–1546.
- J Edward Hu, Huda Khayrallah, Ryan Culkin, Patrick Xia, Tongfei Chen, Matt Post, and Benjamin Van Durme. 2019. Improved lexically constrained decoding for translation and monolingual rewriting. 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 839–850.
- Shuyang Jiang, Yuhao Wang, and Yu Wang. 2023. Selfevolve: A code evolution framework via large language models. *CoRR*, abs/2306.02907.
- Huayang Li, Guoping Huang, Deng Cai, and Lemao Liu. 2020. Neural machine translation with noisy lexical constraints. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 28:1864–1874.
- Aman Madaan, Niket Tandon, Prakhar Gupta, Skyler Hallinan, Luyu Gao, Sarah Wiegrefe, Uri Alon, Nouha Dziri, Shrimai Prabhunoye, Yiming Yang, Sean Welleck, Bodhisattwa Prasad Majumder, Shashank Gupta, Amir Yazdanbakhsh, and Peter Clark. 2023. Self-refine: Iterative refinement with self-feedback. *CoRR*, abs/2303.17651.
- Ian R McKenzie, Alexander Lyzhov, Michael Pieler, Alicia Parrish, Aaron Mueller, Ameya Prabhu, Euan McLean, Aaron Kirtland, Alexis Ross, Alisa Liu, et al. 2023. Inverse scaling: When bigger isn’t better. *arXiv preprint arXiv:2306.09479*.
- Alexander Molchanov, Vladislav Kovalenko, and Fedor Bykov. 2021. Prompt systems for wmt21 terminology translation task. In *Proceedings of the Sixth Conference on Machine Translation*, pages 835–841.
- Yasmin Moslem, Rejwanul Haque, John D. Kelleher, and Andy Way. 2023. Adaptive machine translation with large language models. In Mary Nurminen, Judith Brenner, Maarit Koponen, Sirkku Latomaa, Mikhail Mikhailov, Frederike Schierl, Tharindu Ranasinghe, Eva Vanmassenhove, Sergi Alvarez Vidal, Nora Aranberri, Mara Nunziatini, Carla Parra Escartín, Mikel L. Forcada, Maja Popovic, Carolina Scarton, and Helena Moniz, editors, *Proceedings of the 24th Annual Conference of the European Association for Machine Translation, EAMT 2023, Tampere, Finland, 12-15 June 2023*, pages 227–237. European Association for Machine Translation.
- Yongyu Mu, Abudurexiti Rehemani, Zhiquan Cao, Yuchun Fan, Bei Li, Yinqiao Li, Tong Xiao, Chunliang Zhang, and Jingbo Zhu. 2023. Augmenting large language model translators via translation memories. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 10287–10299.
- 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.
- Silviu Pitis, Michael R. Zhang, Andrew Wang, and Jimmy Ba. 2023. Boosted prompt ensembles for large language models. *CoRR*, abs/2304.05970.

- Matt Post and David Vilar. 2018. Fast lexically constrained decoding with dynamic beam allocation for neural machine translation. In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 1314–1324.
- Matt Post. 2018. A call for clarity in reporting bleu scores. In *Proceedings of the Third Conference on Machine Translation: Research Papers*, pages 186–191.
- Ricardo Rei, Craig Stewart, Ana C Farinha, and Alon Lavie. 2020. Comet: A neural framework for mt evaluation. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*, pages 2685–2702.
- Dimitar Shterionov, Félix do Carmo, Joss Moorkens, Murhaf Hossari, Joachim Wagner, Eric Paquin, Dag Schmidtke, Declan Groves, and Andy Way. 2020. A roadmap to neural automatic post-editing: an empirical approach. *Machine Translation*, 34:67–96.
- Kai Song, Yue Zhang, Heng Yu, Weihua Luo, Kun Wang, and Min Zhang. 2019. Code-switching for enhancing nmt with pre-specified translation. In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 449–459.
- Chau Tran, Shruti Bhosale, James Cross, Philipp Koehn, Sergey Edunov, and Angela Fan. 2021. Facebook ai wmt21 news translation task submission. In *Proceedings of the Sixth Conference on Machine Translation*, pages 205–215.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Advances in Neural Information Processing Systems 30: Annual Conference on Neural Information Processing Systems 2017*, pages 5998–6008.
- Elena Voita, Rico Sennrich, and Ivan Titov. 2019. Context-aware monolingual repair for neural machine translation. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing*, pages 877–886.
- Thuy-Trang Vu and Reza Haffari. 2018. Automatic post-editing of machine translation: A neural programmer-interpreter approach. In *Empirical Methods in Natural Language Processing 2018*, pages 3048–3053. Association for Computational Linguistics (ACL).
- Yuguang Wang, Shanbo Cheng, Liyang Jiang, Jiajun Yang, Wei Chen, Muze Li, Lin Shi, Yanfeng Wang, and Hongtao Yang. 2017. Sogou neural machine translation systems for wmt17. In *Proceedings of the Second Conference on Machine Translation*, pages 410–415.
- Ke Wang, Shuqin Gu, Boxing Chen, Yu Zhao, Weihua Luo, and Yuqi Zhang. 2021a. Termmind: Alibaba’s wmt21 machine translation using terminologies task submission. In *Proceedings of the Sixth Conference on Machine Translation*, pages 851–856.
- Weixuan Wang, Wei Peng, Xupeng Meng, and Qun Liu. 2021b. Huawei aarc’s submissions to the wmt21 biomedical translation task: Domain adaption from a practical perspective. In *Proceedings of the Sixth Conference on Machine Translation*, pages 868–873.
- Shuo Wang, Peng Li, Zhixing Tan, Zhaopeng Tu, Maosong Sun, and Yang Liu. 2022a. A template-based method for constrained neural machine translation. In *Proceedings of the 2022 Conference on Empirical Methods in Natural Language Processing*, pages 3665–3679.
- Shuo Wang, Zhixing Tan, and Yang Liu. 2022b. Integrating vectorized lexical constraints for neural machine translation. In *Proceedings of the 60th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 7063–7073.
- Jinghui Yan, Jiajun Zhang, JinAn Xu, and Chengqing Zong. 2019. The impact of named entity translation for neural machine translation. In *Machine Translation: 14th China Workshop, CWM T 2018, Wuyishan, China, October 25-26, 2018, Proceedings 14*, pages 63–73. Springer.
- Zixin Zeng, Rui Wang, Yichong Leng, Junliang Guo, Xu Tan, Tao Qin, and Tie-yan Liu. 2023. Extract and attend: Improving entity translation in neural machine translation. In *Findings of the Association for Computational Linguistics: ACL 2023*, pages 1697–1710.
- Jiacheng Zhang, Huanbo Luan, Maosong Sun, Feifei Zhai, Jingfang Xu, and Yang Liu. 2021. Neural machine translation with explicit phrase alignment. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 29:1001–1010.

	CCR%	Cost(s)	CCR%	Cost(s)
Direction	English-Chinese		English-Russian	
Iteration0	92.6	1.48	87.9	2.28
Iteration1	93.9	2.21	93.0	3.26
Iteration2	94.6	2.02	94.8	2.99
Iteration3	95.0	2.39	95.1	3.17
Iteration4	95.3	2.56	96.2	3.02
Iteration5	95.6	2.42	96.9	3.25

Table 9: The processing time for each data point at various stages using the Qwen-14b-chat model on an A100 GPU.

	CCR%	Cost(\$)	CCR%	Cost(\$)
Direction	English-Chinese		English-Russian	
Iteration0	92.6	0.62	95.8	0.75
Iteration1	95.1	1.20	97.0	1.28
Iteration2	95.6	1.26	97.3	1.08
Iteration3	95.9	1.31	97.5	1.23
Iteration4	95.9	1.34	97.6	1.23
Iteration5	95.9	1.34	97.6	1.21

Table 10: The monetary costs involved in processing every 1,000 data points using gpt3.5-turbo-0613 at different stages.

Chenrui Zhang, Lin Liu, Jinpeng Wang, Chuyuan Wang, Xiao Sun, Hongyu Wang, and Mingchen Cai. 2023a. PREFER: prompt ensemble learning via feedback-reflect-refine. *CoRR*, abs/2308.12033.

Huaao Zhang, Qiang Wang, Bo Qin, Zelin Shi, Haibo Wang, and Ming Chen. 2023b. Understanding and improving the robustness of terminology constraints in neural machine translation. In *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 6029–6042.

Kechi Zhang, Zhuo Li, Jia Li, Ge Li, and Zhi Jin. 2023c. Self-edit: Fault-aware code editor for code generation. In Anna Rogers, Jordan L. Boyd-Graber, and Naoaki Okazaki, editors, *Proceedings of the 61st Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, ACL 2023, Toronto, Canada, July 9-14, 2023, pages 769–787. Association for Computational Linguistics.

A Cost of Iterations

After undergoing revisions, our TAR significantly enhances the performance of translation systems in constrained translation scenarios. However, multiple rounds of iteration introduce additional computational and financial costs. To quantitatively assess these extra expenditures, we conducted evaluations on the WMT21 terminology translation dataset. Results of time cost are shown in Table 9, and the monetary costs are shown in Table 10. We repeated these tests 10 times for each stage and reported the average scores.

Our results indicate a trend of diminishing returns beyond the third iteration, both in terms of performance and cost-efficiency. Specifically, the time cost stabilizes at 10 seconds per data point and the monetary expense at approximately \$4.5 per 1,000 data points by this iteration. Considering that TAR can significantly enhance the CCR in constrained translations, we believe that the cost is completely within an acceptable range.

B Impact of the number of constraints

In conventional constrained translation methods, it has been observed that as the number of constraints within a single sentence increases, the CCR shows a decreasing trend. To investigate whether our method encounters the same challenge, we conducted experiments on the RTT (Zhang et al., 2023b) dataset. This dataset is composed of 500 samples meticulously selected by linguistic experts from the WMT 13-18 English-German translation test sets, with each sample being accompanied by at least six constraints. These constraints were chosen from a carefully curated set of noun phrases (e.g., names of organizations, individuals, movies, and brands) and common expressions.

To simulate different numbers of constraints, we adopted a method similar to that described in (Zhang et al., 2023b), assuming each sentence in the test set corresponds to N constraints, we randomly selected




Figure 5: (a) The impact of increasing the number of constraints on BLEU. (b) The effect of increasing the number of constraints on CCR.

Wiktionary	BLEU	COMET	CCR %
LLM Const. Trans.	32.1	87.3	97.2
+Revision	32.0	87.3	98.9
WMT21 TT en-zh	BLEU	COMET	CCR %
LLM Const. Trans.	36.4	86.9	92.6
+Revision	35.9	86.9	95.9
ETC en-ru	BLEU	COMET	CCR %
LLM Const. Trans.	32.6	89.9	88.8
+Revision	32.5	89.8	89.8

Table 11: Comparison of metrics before and after revision on the Wiktionary, WMT21 TT en-zh, and ETC en-ru datasets.

Constraints	virus spread → 病毒传播 & Wuhan → 武汉 or 武汉市
Source	In early and mid-January 2020, the virus spread to other Chinese provinces, helped by the Chinese New Year migration and Wuhan being a transport hub and major rail interchange.
Reference	在2020年1月初至1月中旬, 受中国春节人口大流动和武汉作为交通枢纽和主要铁路枢纽的影响, 病毒传播到了中国其他省份。
Const. Trans.	2020年1月初和月中旬, 病毒通过中国春节迁徙和武汉作为交通枢纽和主要铁路换乘站的帮助, 传播到其他中国省份。
+Revision	2020年1月初和月中旬, 病毒传播到其他中国省份, 得益于中国春节迁徙和武汉市作为交通枢纽和主要铁路换乘站的地位。

Table 12: A case demonstrated that during the revision phase, the constraint of translating “virus spread” to “病毒传播” was completed. Meanwhile, the model reorganized the sentence structure and adjusted the wording, resulting in a 5.25 decrease in BLEU score. However, the overall fluency did not change.

between 1 to N constraints for testing. Consequently, we constructed k test subsets, with the number of constraints ranging from 1 to k , Figure 5 presents the results for two metrics ($BLEU, CCR$) as the number of constraints ($k = 6$) varies. From the results, we can observe:

(1) As the number of constraints increases, the BLEU score of a standard LLM does not change, but the BLEU of both our proposed constraint-aware translator and TAR show a certain degree of improvement. When the number of constraints is 6, the BLEU of both methods can increase by about 5.5 points. This phenomenon is understandable because providing more constraints also means that more key parts of the translation sentence are already perceived by the model on how they should be translated.

(2) Similar to traditional constrained translation methods, using LLM as a translator alone, its CCR decreases as the number of constraints increases. However, the CCR of our proposed method does not change significantly, which also indicates that TAR can handle a greater number of constrained translation scenarios in real-world contexts.

C Analysis of Causes for BLEU Score Decline During the Revision Stage

We observed that, among all twelve language directions, three exhibited a decline in BLEU scores after the revision stage. Theoretically, completing more constraints correctly during revision should lead to an increase in BLEU, a string-matching-based metric. However, in reality, to fulfill specific constraints, the model may employ different words or rephrase entire sentences to ensure semantic coherence, which can result in a decrease in BLEU scores. To investigate whether the revision stage could potentially degrade translation quality, we measured the changes in COMET scores before and after revision for these three language directions. The results, as shown in Table 11, indicate that although BLEU scores declined, the COMET scores remained stable, suggesting that TAR does not compromise overall translation quality. We provide a case study in Table 12 to further illustrate this phenomenon.

D Applying mainstream constraint translation methods to LLMs

In the NMT era, many studies have explored training NMT models with data augmentation to develop constrained translation capabilities. Representative approaches include Code-switching, Append, etc. To investigate the effectiveness of these methods on LLMs, we performed experiments using the IATE and

Method	BLEU	CCR%	BLEU	CCR%
<i>Dataset</i>	<i>IATE</i>		<i>Wiktionary</i>	
LLM Trans	32.0	85.4	32.0	88.7
LLM Code-switching	31.8	94.0	31.9	92.6
LLM Append	31.6	93.3	31.5	91.9
LLM Const. Trans.	32.0	96.2	32.1	97.2
+Revision	32.0	98.9	32.0	98.9

Table 13: Performance of traditional data augmentation methods on LLMs.

Wiktionary datasets. The results are presented in Table 13. Although code-switching prompts are somewhat effective, they typically decrease translation quality and worsen CCR metrics compared to natural language prompts. We speculate that this is because LLMs have not encountered similar prompt formats during training, leading to alignment issues during task execution. Adopting techniques such as few-shot learning and fine-tuning may mitigate issues of misalignment between prompts and model training data. However, our proposed TAR method uses natural language prompts, which are more suitable for the processing style of large language models, thus enabling better understanding and execution of translation tasks.

CCL 2024