

CONTROL PREFIXES for Parameter-Efficient Text Generation

Jordan Clive
Imperial College London

Kris Cao
DeepMind, London, UK

Marek Rei
Imperial College London

{jordan.clive19, marek.rei}@imperial.ac.uk
kriscao@deepmind.com

Abstract

Prefix-tuning is a powerful lightweight technique for adapting a large pre-trained language model to a downstream application. However, it uses the same dataset-level tuned prompt for all examples in the dataset. We extend this idea and propose a dynamic method, CONTROL PREFIXES, which allows for the inclusion of conditional input-dependent information, combining the benefits of prompt tuning and controlled generation. The method incorporates attribute-level learnable representations into different layers of a pre-trained transformer, allowing for the generated text to be guided in a particular direction. We provide a systematic evaluation of the technique and apply it to five datasets from the GEM benchmark for natural language generation (NLG). Although the aim is to develop a parameter-efficient model, using only 0.1–3% trainable parameters, we show CONTROL PREFIXES can even outperform full fine-tuning methods. We present state-of-the-art results on several data-to-text datasets, including WebNLG.

1 Introduction

Recently, approaches in text generation have been dominated by adapting one large-scale, pre-trained language model (PLM) to various downstream tasks. Such adaptation is often performed via fine-tuning, which necessitates updating and storing all of the parameters, resulting in multiple new language models (LMs), one for each task. This poses a considerable challenge to the deployment of NLP systems in practice, especially as the scale of PLMs continues to climb from millions to billions of parameters. Moreover, full fine-tuning has been shown to be unnecessarily profligate through overwriting natural language understanding (NLU) that could otherwise be shared among tasks (Peters et al., 2019); it has also been shown that fine-tuned networks do not deviate substantially from the pre-trained one in parameter space (Aghajanyan et al.,

2020; Radiya-Dixit and Wang, 2020), implying the existence of parameter efficient alternatives.

Many researchers have sought to alleviate these issues by using *fixed-LM* techniques, where the parameters of the base LM remain unchanged. An ever-growing subset of these methods can be considered prompt tuning, where language models are adapted to downstream tasks with the aid of a tuned prompt accompanying the input. A recent survey on prompt tuning (Liu et al., 2021a), however, notes the dearth of research exploring *dynamic* prompts, which are input-dependent. This work fills this gap in the literature and considers such dynamic prompts. Existing controlled generation techniques either aim to generate text with specific target qualities, independent of overall task performance, or are methods that have the benefit of updating not only the attribute-level parameters but training all the parameters in the language model.

We propose the *dynamic* prompting method CONTROL PREFIXES. The method extends prefix-tuning and integrates static task-specific prompts at every layer of a model, adding only 0.1–3% additional parameters to the base LM. With CONTROL PREFIXES we aim to preserve the *fixed-LM* property, while also allowing datapoint-specific attributes to act as guidance signals at the input-level. This is done by employing modular *control prefixes*, which change alongside the input according to the guidance signal. Operating together with the static prompt parameters, these dynamic prompts can steer the frozen PLM to extend finer-grained control. The chosen attributes can provide additional information about the input, for example the domain of a data-to-text triple set, or it can specify some aspect of the desired output, such as the target length for text simplification.

We evaluate our method on an array of text generation tasks, leveraging additional input-level information specific to each dataset. Our results show that our parameter efficient architecture out-

performs previous approaches, many of them based on full fine-tuning, according to the WebNLG (Gardent et al., 2017), DART (Radev et al., 2020) and E2E Clean (Dušek et al., 2019) data-to-text datasets. In addition, our method attains higher human-assessed performance than existing systems for summarization on XSum (Narayan et al., 2018). Although CONTROL PREFIXES no longer operates in the standard setting for NLG tasks, by being not confined to just using the textual input, we focus on datasets where the attribute-level information is available as part of the task.

We also consider the common case where the attribute-level information is not available, and demonstrate that zero-shot learning with CONTROL PREFIXES can be effective. We show similar control prefix representations are learned by the model for semantically similar attribute labels.

2 Related Work

Prompt Tuning Unlike the discrete text prompts used by GPT-3 (Brown et al., 2020), in prompt tuning, soft prompts are learned through back-propagation to maximize the information from labelled data. This work focuses on tuning methods as zero-shot prompting performance lags far behind tuned models on supervised datasets (Lester et al., 2021). Several successive works (Logeswaran et al., 2020; Liu et al., 2021b; Lester et al., 2021) employ prompt-embedding tuning, which trains continuous embeddings prepended to the input embeddings. Li and Liang (2021) discovered that prefix-tuning was more effective than prompt-embedding tuning for text generation. In prefix-tuning, additional trainable key-value pairs, which are fixed across all examples, are used to augment the left context in every attention computation. Therefore, the prompt has constituents at every layer rather than being confined to steer the frozen LM only through the input as in embedding tuning.

Controlled generation A complementary field to prompt learning is controlled generation, which aims to incorporate various types of guidance (e.g. length specifications (Kikuchi et al., 2016) or highlighted phrases (Grangier and Auli, 2018)) beyond the input text into the generation model. Johnson et al. (2016) successfully trained a multilingual translation model with control tokens to encode each language. Keskar et al. (2019) pre-trained a 1.63B parameter model, also alongside conditional control tokens, and demonstrated these learnt to

govern style, content, and task-specific behaviour. However, these models require the whole underlying LM to be fine-tuned alongside the control tokens for a particular task.

Alternatives exist, such as plug-and-play perturbations of the LM hidden states towards a target attribute (Nguyen et al., 2016; Dathathri et al., 2020). These methods use fixed LMs and are able to control target qualities such as sentiment and topic. However, they are slow at inference time due to requiring multiple passes for a single batch. The shift in conditional probability has also been shown to increase text degeneration (Holtzman et al., 2019).

Dynamic prompts There have been few works exploring *dynamic* prompts (Liu et al., 2021a; Tsimpoukelli et al., 2021), which are input-dependent. Perhaps most similar to our work is work by Yu et al. (2021), who use an attribute alignment function to form dynamic prompts. Unlike our work, the prompt does not have a static component and aims to generate text with specific target attributes, independent of task performance. With CONTROL PREFIXES, the intention is to also maximize task-specific performance, which is why we maintain a large static prompt component to specify the task itself.

3 CONTROL PREFIXES

3.1 Background

This work considers sequence-to-sequence tasks where the objective is to model the conditional probability $P(Y | X)$ with X and Y representing the tokenized input and output sequences respectively. For example, in summarization, X could be an article and Y would be a short target summary.

In this work we experiment with T5-large (Raffel et al., 2020) and BART_{LARGE} (Lewis et al., 2020) as the underlying pre-trained LMs with parameters ϕ ; and as we consider fixed-LM methods, ϕ **always remains frozen**. These models are Transformer encoder-decoder models where decoding proceeds auto-regressively. Let us denote d to represent the hidden state dimension and L the number of layers. We use (E, Dc, Dm) to denote the three classes of attention present in each layer: self-attention in the encoder (E), decoder cross-attention (Dc) and decoder masked-attention (Dm). For an attention computation in the l -th layer, the query, key and value matrices are denoted $Q_l \in \mathbb{R}^{N \times d}$, and $K_l, V_l \in \mathbb{R}^{M \times d}$, where N is the number of tokens in the series relating

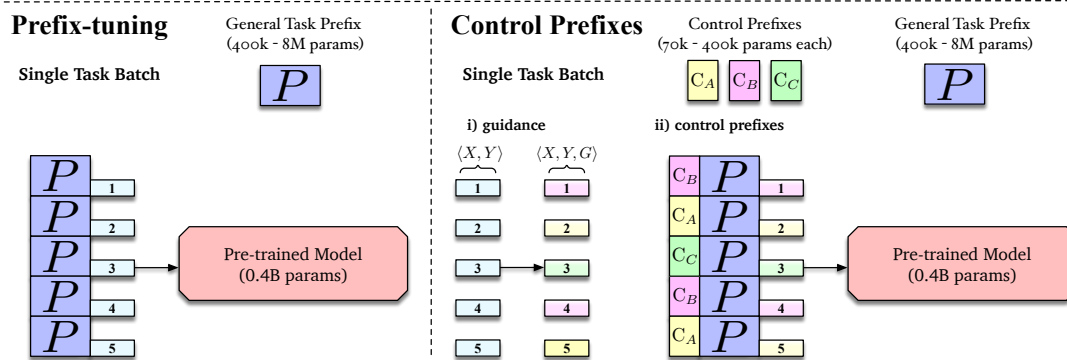


Figure 1: High-level diagram contrasting prefix-tuning and CONTROL PREFIXES in the single-task setup for a PLM such as BART_{LARGE}. The same single-task batch (examples 1,2,3,4 and 5) is considered for both setups. Left: Prefix-tuning has one general prefix P for all examples. Right: CONTROL PREFIXES utilizes additional attribute information at the input-level, G , in **i)**. This conditional information is used in **ii)** to dictate which control prefix (C_A , C_B , C_C) to use for a particular example in a batch. This takes advantage of prefix-tuning’s capacity to include different prefixes in one forward pass.

to queries, and M is the number of tokens in the series relating to keys and values.

3.2 Intuition

Using a fixed PLM that captures broad natural language understanding provides the model with a parameter-efficient starting point which can be shared by many different tasks. Combining this with a trainable task representation allows the model to learn information relevant to one particular task. Furthermore, introducing attribute-level parameters allows us to guide the generation into a required direction and provide the model with datapoint-level information. The general task-specific parameters can themselves adapt to the modular *control prefixes*, which change according to the guidance signal for each input X . This demarcation of parameters enables fine-grained control to be extended to aid performance on downstream tasks. CONTROL PREFIXES can therefore leverage input-level information while being a fixed-LM, parameter efficient method.¹ For this work, we only consider discrete labels as attributes for the guidance signal.

3.3 Description

The model uses a general task prefix P_θ ("task-specific parameters") and also trains a set of control prefixes C_θ that change depending on the input ("attribute-level parameters"). This requires attribute-level information or guidance G , to indicate which control prefixes to be used while pro-

cessing a given input X .² Let us consider the parallel corpus $\mathcal{Z} = \{\langle X^j, Y^j, G^j \rangle\}_{j=1, \dots, N}$, where G^j indicates all the conditional attribute-level information for the sample j . The goal is to optimize through gradient descent the final inference parameters, θ , whilst the underlying ϕ parameters of the pre-trained LM remain frozen:

$$\theta^* = \arg \max_{\theta} \sum_{j=1}^N \log p(Y^j | X^j, G^j; P_\theta, C_\theta, \phi). \quad (1)$$

General Prefix For each attention class (E, Dc, Dm), a distinct prefix of key-value pairs is learnt, $P = \{P_1, \dots, P_L\}$, where $P_l \in \mathbb{R}^{\rho \times 2d} \forall l \in \{1, \dots, L\}$. $P \in \mathbb{R}^{\rho \times 2dL}$ and ρ is the prompt length, i.e. the number of additional key-value pairs in each attention computation. In prefix-tuning³, for an attention computation in the l -th layer, K_l and V_l are augmented to become

$$K'_l = [P_{l,K}; K_l], V'_l = [P_{l,V}; V_l] \quad (2)$$

where $K'_l, V'_l \in \mathbb{R}^{(\rho+M) \times d}$. The overall general prefix, parameterized by θ , is $P_\theta = \{P^E, P^{Dc}, P^{Dm}\}$, where $P_\theta \in \mathbb{R}^{\rho \times 6dL}$.

Control Prefixes Let us consider one attribute with R possible labels⁴, such as the news domain of an article (e.g. sport, technology etc.),

²We discuss cases where G is not present in §6.2.

³There has been confusion in recent work concerning different forms of prefix-tuning (Li and Liang, 2021). For details and observations of the benefits (previously unremarked upon) conferred by key-value pair prefix-tuning, see Appendix C.

⁴The procedure can be generalized to multiple attributes; we use up to four attributes and varying control prompt lengths.

¹We use the term parameter efficient to denote methods adding <3% additional parameters to a fixed LM’s parameters.

$C_\theta = \{C_{\theta,1}, \dots, C_{\theta,R}\}$, where $C_{\theta,r} \in \mathbb{R}^{\rho_c \times 6dL}$, $\forall r \in \{1 \dots R\}$. $C_{\theta,r}$ represents the control prefix learnt for the r -th attribute label and the parameter ρ_c denotes the control prompt length for this *particular* attribute. Let \mathcal{A} be a function which returns the corresponding control prefix for the attribute label indicated by G . In CONTROL PREFIXES the K_l and V_l are augmented to become

$$\begin{aligned} K_l'' &= [\mathcal{A}(G)_{l,K}; P_{l,K}; K_l], \\ V_l'' &= [\mathcal{A}(G)_{l,V}; P_{l,V}; V_l] \end{aligned} \quad (3)$$

where $K_l'', V_l'' \in \mathbb{R}^{(\rho_c + \rho + M) \times d}$.

Shared Re-parameterization Li and Liang (2021) found that prefix optimization is stabilized by increasing the number of trainable parameters. This is achieved by introducing a feed-forward network to re-parameterize the prefix. Rather than one network, we use three distinct two-layered large feed-forward neural networks for each attention class, applied row-wise. For each attention class (E, Dc, Dm), $P = \text{MLP}(\tilde{P})$ where $\tilde{P} \in \mathbb{R}^{\rho \times d}$ is smaller than the matrix $P \in \mathbb{R}^{\rho \times 2dL}$, and each MLP has an intermediate dimension k which we set to 800. The distinct MLPs and each \tilde{P} are parameterized by training parameters $\tilde{\theta}$; thus, θ is a function of $\tilde{\theta}$ and $|\theta| < |\tilde{\theta}|$. Once training is complete, the final θ parameters can be saved for use at inference and the re-parameterization parameters dispensed with.

As described for the general prefix, P_θ , each control prefix, $C_{\theta,r}$, comprises three constituents for each attention class: $C_{\theta,r} = \{C_r^E, C_r^{Dc}, C_r^{Dm}\}$. The re-parameterization of $C_{\theta,r}$ occurs in the same manner as P_θ , sharing the same MLP^E , MLP^{Dc} and MLP^{Dm} . When using a disjoint set of re-parameterizations for the control prefixes, learning becomes unstable and performance degrades.⁵

Recent work by Buhai et al. (2020) show that over-parameterization can smooth the optimization landscape. With this in mind, the three distinct re-parameterizations compel each prefix element to coordinate control for the particular attention class. For example, the rows of P^E and C_r^E lie in a vector space better coordinated for moderating the processing of the input sequence X than P^{Dm} and C_r^{Dm} . This is due to being formed from the shared mapping MLP^E .

⁵This would also result in a significant increase in the number of training parameters $\tilde{\theta}$. In contrast, with the methodology outlined, each additional control prefix relates to only an additional $d\rho_c$ training parameters.

4 Experimental Setup

4.1 Datasets, Guidance and Metrics

Examples of specific attribute labels for each task are found in the Appendix.⁶

Data-to-text The objective of data-to-text generation is to produce fluent text from structured input, such as a triple set (a set of subject-predicate-objects). Following Li and Liang (2021), we evaluate on the data-to-text datasets DART (Radev et al., 2020) and WebNLG (Gardent et al., 2017). However, we implement prefix-tuning for T5-large rather than GPT-2, as T5-large provides a stronger baseline and enables comparison with state-of-the-art (SOTA) systems.⁷ We also report results on E2E Clean (Dušek et al., 2019), a dataset focused on the restaurant domain. We use the official evaluation scripts and report BLEU (Papineni et al., 2002), METEOR (Lavie and Agarwal, 2007), and TER (Snoover et al., 2006) metrics.⁸

WebNLG contains triple sets from DBpedia (Auer et al., 2007). The test set is divided into two partitions: ‘‘Seen’’, which contains 10 DBpedia categories present in the training set, and ‘‘Unseen’’, which covers 5 categories never seen during training.⁹ These categories, such as *Airport* or *Food* are used as a guidance signal in our experiments (indicated by A_1 in Table 1); our approach for unseen categories is discussed in §6.2.

Providing the category explicitly as guidance with CONTROL PREFIXES may enable properties of triples belonging to a specific WebNLG category to be captured more effectively. This intuition is supported by studies showing a clear disparity in the performance of different model types between different categories (Moryossef et al., 2019; Castro Ferreira et al., 2020). DART is an open-domain, multi-source corpus, with six sources: internal and external human annotation of both Wikipedia tables and WikiSQL, as well as the two existing datasets WebNLG and E2E Clean. Radev et al. (2020) showed fine-tuning T5-large on the WebNLG dataset with only the human an-

⁶For data-to-text see Tables 13, 14; summarization see Table 15 and simplification see Table 11.

⁷BART_{LARGE} exhibits inferior performance to T5-large on data-to-text; for example, 9.7 BLEU points lower on WebNLG Unseen (Ribeiro et al., 2020).

⁸Full results from the evaluation scripts, including machine-learned metrics can be found in Appendix A.

⁹All the training category labels are visible in Appendix B, where we visualize control prefixes corresponding to each training category.

notated portion of DART achieves SOTA performance, whilst using the whole DART dataset is not as effective. Nevertheless, this inspired the idea of using the six DART sub-dataset sources as a controllable attribute, represented by A_2 in Table 1. This strategy was inspired by previous work which incorporates auxiliary scaffold tasks (Swayamdipta et al., 2018; Cohan et al., 2019; Cachola et al., 2020).

Simplification We use WikiLarge (Zhang and Lapata, 2017) as the training data and evaluate on two simplification benchmarks: TurkCorpus (Xu et al., 2016) and ASSET (Alva-Manchego et al., 2020). Both benchmarks are composed of the same 2000 validation source and 359 test source sentences. However, the 10 ASSET references per source focus on a more diverse set of rewriting simplifications than the 8 TurkCorpus references per source. Martin et al. (2020) introduced ‘BART_{LARGE} with ACCESS’, which is a fine-tuned BART_{LARGE} model trained alongside control tokens to condition on four simplification-specific attributes, such as the length compression ratio (the length of the target sequence relative to the source sequence). We use the same controllable attributes in this work to directly compare with Martin et al. (2020) (Table 2). The control ratios are discretized into bins of fixed-width 0.05, capped to a maximum ratio of 2. At inference time, once the model has been trained with these oracle controls, the control ratios are set to desired values by tuning on the respective validation set.

We report the non-learned metrics SARI (Xu et al., 2016) and FKGL (Kincaid et al., 1975).¹⁰ Unlike previous studies, we also use the machine-learned Q&A metric QuestEval (Scialom et al., 2021) to assess our text simplification models.

Summarization As in Li and Liang (2021), we report results on the XSum dataset (Narayan et al., 2018) using BART_{LARGE}. XSum comprises 226,711 British Broadcasting Corporation (BBC) articles coupled with their single-sentence summaries, where each sample corresponds to a unique URL. The URL contains information on whether the sub-directory is from the BBC Sport or BBC News page (A_1 in Table 3), and further sub-directory information (A_2 in Table 3, where A_2 has 40 labels), for example (‘sport’, ‘formula1’) or

¹⁰We use the FKGL and the latest version of SARI implemented in EASSE (Alva-Manchego et al., 2019) which is used in Martin et al. (2020).

(‘news’, ‘science’). The motivation for using this as guidance is that different sub-directories are likely to share properties relating to how the information is presented; journalists are also usually confined to one domain. We report on the customary ROUGE scores (Lin, 2004).

4.2 Training Details

For the data-to-text datasets, we follow Ribeiro et al. (2020) and linearize the triples, prepending the special tokens <H>, <R>, and <T> before the subject, predicate, and object of an individual triple.¹¹ We also prepend “translate Graph to English: ” to every input (Raffel et al., 2020). Full training and hyperparameter details can be found in Appendix D.

5 Results

5.1 Data-to-Text

Results in Table 1 show that for DART, both CONTROL PREFIXES (A_2) and prefix-tuning attain higher performance than the current SOTA, which is T5-large fine-tuned (Radev et al., 2020), by 1.29 and 0.54 BLEU points respectively. This indicates CONTROL PREFIXES can exert control over the frozen T5-large more effectively than prefix-tuning.

The SOTA for WebNLG is a T5-large model fine-tuned on WebNLG and the human annotated portion of DART (Radev et al., 2020).¹² Compared to this model, CONTROL PREFIXES achieves a 0.83 higher BLEU overall, and 1.33 on the Seen categories. Notably, CONTROL PREFIXES (A_1) outperforms CONTROL PREFIXES (A_1, A_2) on the Seen component of the dataset, but does not generalize as well to the unseen categories. We argue that this illustrates the benefit of using both controllable attributes. The prefix-tuning model with additional DART data, like the SOTA, is trained on only the human annotated portion and yields a minor performance increase of 0.05 BLEU compared to prefix-tuning solely trained on WebNLG. We believe this indicates that for fine-tuning, training on a complementary type of additional data allows the PLM to maintain more NLU by not over-fitting a narrow distribution, leading to better LM generalization. In contrast, for prefix-tuning, much of

¹¹The embeddings relating to these special tokens are the only embeddings we train, as our work is focused on fixed-LM methods.

¹²Additional training data is permitted by the organizers of the E2E Clean and WebNLG datasets.

	$\phi\%$	DART			$\phi\%$	WebNLG			$\phi\%$	E2E Clean	
		BLEU	METEOR	TER ↓		S	U	A		BLEU	METEOR
T5-large fine-tuned	100	50.66	40	43	100	64.89	54.01	59.95	100	41.83	38.1
SOTA	100	50.66	40	43	100	65.82	56.01	61.44	100	43.6	39
Prefix-tuning	1.0	51.20	40.62	43.13	1.0	66.95	55.39	61.73	1.0	43.66	39.0
CONTROL PREFIXES (A_1)	-	-	-	-	1.4	67.32	55.38	61.94	-	-	-
+Data: DART											
Prefix-tuning	1.0	51.20	40.62	43.13	1.0	67.05	55.37	61.78	1.0	43.04	38.7
CONTROL PREFIXES (A_2)	1.1	51.95	41.07	42.75	1.0	66.99	55.56	61.83	1.0	44.15	39.2
CONTROL PREFIXES (A_1, A_2)	-	-	-	-	1.4	67.15	56.41	62.27	-	-	-

Table 1: Data-to-text test set results reported on the respective official evaluation scripts. $\phi\%$ denotes the % of additional parameters to the number of fixed-LM parameters required at inference time. T5-large fine-tuned results for WebNLG are from Ribeiro et al. (2020) and for DART are from Radev et al. (2020). Note the results in the main body of the GEM paper (Gehrmann et al., 2021) are reported on the validation set, rather than the test set as is done here. Several of the baseline results were only reported to the significant figures shown. A_1 signifies models trained with control prefixes for the *WebNLG category* attribute, and A_2 with control prefixes for the DART *sub-dataset source* attribute. For WebNLG, S, U and A refer to BLEU scores for the *Seen*, *Unseen* and *All* portions of the dataset. The DART results are reported on the official evaluation script for v1.1.1, the same version as the official leaderboard. A CONTROL PREFIXES model attains state-of-the-art results for each dataset.

this gain has already been realized by retaining the original frozen parameters.

The SOTA (Harkous et al., 2020) for E2E Clean consists of a fine-tuned GPT-2 with a semantic fidelity classifier trained on additional generated data. CONTROL PREFIXES (A_2), which can leverage the heterogeneous DART datasets, outperforms this model in terms of the BLEU score.

5.2 Simplification

Table 2 reveals that prefix-tuning BART performs comparably to fine-tuning BART. When comparing our CONTROL PREFIXES to fine-tuned ‘BART_{LARGE} with ACCESS’ there is comparable performance in terms of SARI for ASSET, and better FKGL results on ASSET. For text simplification, Martin et al. (2020) indicate the gains from using the controllable attributes, as assessed by SARI and FKGL, are mostly due to being able to calibrate the length ratio, with validation and test sets being drawn from the same distribution, as opposed to the WikiLarge training distribution. CONTROL PREFIXES also achieves higher SARI and FKGL scores on TurkCorpus compared to the *Gold Reference*, which evaluates against other human annotators.

5.3 Summarization

There is considerable inconsistency regarding author-conducted human evaluation for NLG (van der Lee et al., 2021). Therefore, we opted to submit our CONTROL PREFIXES model outputs to an externally run evaluation framework, GENIE

(Khashabi et al., 2021), which provides an unbiased attestation of performance. Their sample size of 300 examples is larger than the 50 or 100 examples that have been previously used for XSum and is typical of human evaluation experiments (Narayan et al., 2018; Dou et al., 2020). Both human evaluation and automated ROUGE metrics can be seen in Table 3. The confidence intervals indicate that this result is not necessarily definitive, but it also highlights that the quality of generations in this domain is not captured fully by ROUGE. For the datasets considered, the automatic metrics are the least reliable for XSum as it is the only dataset with a single gold reference.

The results also show that CONTROL PREFIXES performs better than prefix-tuning in terms of ROUGE. We are not able to report the same human-assessment results for prefix-tuning, as each participant of GENIE is limited to one submission and there is no existing result for prefix-tuning.

6 Analysis

6.1 Visualizing Control Prefixes

Fig. 2 displays t-SNE (Maaten and Hinton, 2008) visualizations of the length compression control prefixes learnt as part of our simplification CONTROL PREFIXES model.¹³ We plot only the decoder self-attention constituent of each control prefix (comprising multiple key-value pairs at each layer) as the length ratio directly concerns the tar-

¹³A perplexity of 5 is used for all plots.

	$\phi\%$	ASSET			TurkCorpus		
		SARI	FKGL ↓	QuestEval	SARI	FKGL ↓	QuestEval
Gold Reference	-	44.87	6.49	0.63*	40.04	8.77	0.66*
BART _{LARGE} with ACCESS [†]	100	43.63	6.25	0.64*	42.62	6.98	0.66*
BART _{LARGE} fine-tuned	100	39.91*	7.73*	-	39.55*	7.73*	-
Prefix-tuning	1.8	40.12	7.28	-	39.06	7.28	-
CONTROL PREFIXES	1.8	43.58	5.97	0.64	42.32	7.74	0.66

Table 2: Simplification results on ASSET and TurkCorpus test sets. [†]This model is from Martin et al. (2020), where the authors fine-tuned BART_{LARGE} model alongside control tokens for the four attributes. The CONTROL PREFIXES model is trained with control prefixes for these same four attributes. Prefix-tuning and CONTROL PREFIXES use BART_{LARGE} as the fixed LM. The * denotes baseline results calculated in this study—the model outputs of Martin et al. (2020) are publicly available. The BART_{LARGE} with ACCESS and CONTROL PREFIXES model are the average test set results over 5 random seeds. We bold the best results of parameter-efficient models in the results tables, while fully fine-tuned models and human performance are reported for reference.

	$\phi\%$	Human overall	Human conciseness	Human fluency	Human no-hallucination	Human informativeness	ROUGE		
							R-1	R-2	R-3
BART _{LARGE} fine-tuned	100	0.49 ^{+0.03} _{-0.04}	0.50 ^{+0.03} _{-0.03}	0.50 ^{+0.03} _{-0.03}	0.52 ^{+0.03} _{-0.03}	0.49 ^{+0.03} _{-0.03}	45.14*	22.27*	37.25*
PEGASUS fine-tuned	100	0.49 ^{+0.03} _{-0.03}	0.52 ^{+0.02} _{-0.03}	0.49 ^{+0.03} _{-0.02}	0.49 ^{+0.03} _{-0.03}	0.49 ^{+0.03} _{-0.03}	47.21*	24.56*	39.25*
T5 (11B) fine-tuned	100	0.47 ^{+0.03} _{-0.03}	0.49 ^{+0.02} _{-0.02}	0.50 ^{+0.03} _{-0.03}	0.49 ^{+0.03} _{-0.03}	0.48 ^{+0.03} _{-0.03}	-	-	-
Prefix-tuning	3.0	-	-	-	-	-	43.53	20.66	35.63
CONTROL PREFIXES (A_1, A_2)	2.8	0.51 ^{+0.03} _{-0.03}	0.53 ^{+0.02} _{-0.02}	0.51 ^{+0.03} _{-0.03}	0.53 ^{+0.03} _{-0.03}	0.49 ^{+0.03} _{-0.03}	43.81	20.84	35.81

Table 3: Summarization results on XSum. The human-assessed results are from the GENIE benchmark, where the 95% confidence intervals are computed with bootstrap re-sampling. Note the BART_{LARGE} and PEGASUS fine-tuned results for the human-assessed dimensions are transcribed from Khashabi et al. (2021), whilst the automatic metric results, indicated by *, are from Lewis et al. (2020) and Zhang et al. (2019). Prefix-tuning and CONTROL PREFIXES (A_1, A_2) use BART_{LARGE} as the fixed LM. A_1 refers to the BBC news/sport page attribute and A_2 the further sub-directory attribute. We bold the best results of parameter-efficient models in the results tables for ROUGE, with fully fine-tuned models as reference. The public GENIE leaderboard is available at <https://leaderboard.allenai.org/genie-xsum/>.

get.¹⁴ The relationship learnt by the control prefixes is very manifest, aided by the near uniform distribution of length ratios in the WikiLarge training dataset from 0 to 1.1.

Fig. 2 establishes that for this simplistic attribute, different control prefixes corresponding to similar attribute labels (i.e. varying length ratios for the length attribute) share properties. Interestingly the decoder cross-attention of the control prefix is not as manifest. We believe this is due to BART_{LARGE} being accustomed to the same cross-attention key-value pairs in each layer.

6.2 Zero-shot Learning

We argue that even for more complicated attributes, such as the WebNLG category attribute, if the attribute labels are semantically similar, the respective control prefixes will similarly assist the general




Figure 2: t-SNE visualizations for the decoder self-attention constituent of the simplification model’s length compression control prefixes. Each circle represents a control prefix corresponding to each length ratio (bins of fixed width 0.05, from 0 to 1.1).

¹⁴Plots for the encoder and decoder cross-attention constituents can be seen found in Appendix E.

task-specific prefix and the frozen LM during generation. Previous work has discussed the notion of task similarity (Achille et al., 2019) for prompt learning methods (Lester et al., 2021); however, we argue prefixes concerning different labels of one attribute are more likely to overlap in terms of learnable properties than different tasks or whole datasets.

In the case of WebNLG, where although no examples of the unseen category are present during training, a textual label for the category exists. These labels were available to all competition participants. This gives us some prior on the properties of the unseen categories, which we show is enough to successfully zero-shot transfer with control prefixes. For each WebNLG model with the category attribute, we map each category’s textual label, including for the unseen categories, to a Glove embedding¹⁵ (Pennington et al., 2014). Then for each unseen category, we map to the seen category with the highest cosine similarity in embedding space, and use that control prefix at inference for the corresponding unseen sample. For example, the control prefix for the seen category *SportsTeam* is used for examples relating to the unseen category *Athlete*.¹⁶

Table 4 shows a comparison of using an out-of-vocabulary (OOV) control prefix for each example with an unseen category, and the zero-shot transfer method for both WebNLG datasets¹⁷. The OOV control prefix is trained on a random 2% of the data for each accumulated batch. These results indicate that zero-shot transfer is more promising than a learned OOV representation. The result fundamentally depends on the WebNLG categories, and if similar textual labels pertain to similar triple sets that CONTROL PREFIXES can utilize.

6.3 Discussion

We also investigated a simpler architecture ‘prefix-tuning + control tokens’ which informs the model of the identical guidance signal as in CONTROL PREFIXES, but with trainable control tokens instead of control prefixes. Appendix F reveals that CONTROL PREFIXES consistently outperforms prefix-tuning + control tokens on the data-to-text and summarization datasets, while the results are both com-

¹⁵Glove Common Crawl (840B tokens, 2.2M vocab, cased, 300d vectors).

¹⁶Appendix H displays model output for WebNLG along with the zero-shot procedure.

¹⁷We also report results on WebNLG+ 2020 (Castro Ferreira et al., 2020), the second official WebNLG competition, in Appendix B.

	Unseen Component		BLEU
	# Examples	# Categories	
<i>WebNLG</i>	891	5	
OOV Representation			56.35
Zero-shot			56.41
<i>WebNLG+ 2020</i>	896	3	
OOV Representation			50.02
Zero-shot			50.39

Table 4: A comparison of the performance on the *Unseen* portions for WebNLG test sets, with i) a single OOV Control Prefix used for all samples from unseen categories, or ii) the zero-shot transfer approach outlined, utilizing the available textual labels.

parable to the *Gold References* on simplification datasets. This indicates that CONTROL PREFIXES is a superior parameter-efficient framework in leveraging additional information, whilst maintaining the *fixed-LM* property.

The alternative method is less expressive than CONTROL PREFIXES, by only exerting control through the embeddings rather than through each layer. CONTROL PREFIXES fundamentally depends on the strength of the guidance signal and by adding the constraint of attribute information being available with the dataset the guidance signal is naturally weaker. However, we show that CONTROL PREFIXES is a powerful general method which can utilize this signal to achieve a modest but consistent improvement across an array of tasks.

7 Conclusion

We introduce CONTROL PREFIXES, a parameter-efficient controlled generation technique, which integrates a task-specific prompt alongside dynamic prompts to leverage additional input-level information. The method extends prefix-tuning, enabling the model to have finer-grained control over generated text, and assists in maximizing downstream task performance.

We demonstrate that CONTROL PREFIXES outperforms prefix-tuning and prefix-tuning with embedding level guidance, as well as existing approaches, on an array of natural language generation tasks. Our method attains state-of-the-art results on several data-to-text datasets including WebNLG. This is despite learning <2% additional parameters to the underlying LM parameters (which remain fixed). Additionally, our method holds the highest human evaluation ranking on the external platform GENIE for the summarization dataset XSum.

References

- Alessandro Achille, Michael Lam, Rahul Tewari, Avinash Ravichandran, Subhransu Maji, Charless Fowlkes, Stefano Soatto, and Pietro Perona. 2019. [Task2vec: Task embedding for meta-learning](#). In *2019 IEEE/CVF International Conference on Computer Vision (ICCV)*, pages 6429–6438.
- Armen Aghajanyan, Luke Zettlemoyer, and Sonal Gupta. 2020. [Intrinsic dimensionality explains the effectiveness of language model fine-tuning](#). *CoRR*, abs/2012.13255.
- Fernando Alva-Manchego, Louis Martin, Carolina Scarton, and Lucia Specia. 2019. [Easse: Easier automatic sentence simplification evaluation](#). *arXiv preprint arXiv:1908.04567*.
- Fernando Emilio Alva-Manchego, Louis Martin, Antoine Bordes, Carolina Scarton, Benoît Sagot, and Lucia Specia. 2020. [ASSET: A dataset for tuning and evaluation of sentence simplification models with multiple rewriting transformations](#). *CoRR*, abs/2005.00481.
- Sören Auer, Christian Bizer, Georgi Kobilarov, Jens Lehmann, Richard Cyganiak, and Zachary Ives. 2007. [Dbpedia: A nucleus for a web of open data](#). In *The Semantic Web, pages 722–735*, Berlin, Heidelberg. Springer Berlin Heidelberg.
- 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 *NeurIPS*.
- Rares-Darius Buhai, Yoni Halpern, Yoon Kim, Andrej Risteski, and David Sontag. 2020. [Empirical study of the benefits of overparameterization in learning latent variable models](#).
- Isabel Cachola, Kyle Lo, Arman Cohan, and Daniel S. Weld. 2020. [TLDR: extreme summarization of scientific documents](#). *CoRR*, abs/2004.15011.
- Thiago Castro Ferreira, Claire Gardent, Nikolai Ilinykh, Chris van der Lee, Simon Mille, Diego Moussallem, and Anastasia Shimorina. 2020. [The 2020 bilingual, bi-directional WebNLG+ shared task: Overview and evaluation results \(WebNLG+ 2020\)](#). In *Proceedings of the 3rd International Workshop on Natural Language Generation from the Semantic Web (WebNLG+)*, pages 55–76, Dublin, Ireland (Virtual). Association for Computational Linguistics.
- Arman Cohan, Waleed Ammar, Madeleine van Zuylen, and Field Cady. 2019. [Structural scaffolds for citation intent classification in scientific publications](#). *CoRR*, abs/1904.01608.
- Sumanth Dathathri, Andrea Madotto, Janice Lan, Jane Hung, Eric Frank, Piero Molino, Jason Yosinski, and Rosanne Liu. 2020. [Plug and play language models: A simple approach to controlled text generation](#). In *International Conference on Learning Representations*.
- Zi-Yi Dou, Pengfei Liu, Hiroaki Hayashi, Zhengbao Jiang, and Graham Neubig. 2020. [Gsum: A general framework for guided neural abstractive summarization](#). *CoRR*, abs/2010.08014.
- Ondřej Dušek, David M. Howcroft, and Verena Rieser. 2019. [Semantic noise matters for neural natural language generation](#). In *Proc. of the 12th International Conference on Natural Language Generation*, pages 421–426, Tokyo, Japan. Association for Computational Linguistics.
- Claire Gardent, Anastasia Shimorina, Shashi Narayan, and Laura Perez-Beltrachini. 2017. [The WebNLG challenge: Generating text from RDF data](#). In *Proceedings of the 10th International Conference on Natural Language Generation*, pages 124–133, Santiago de Compostela, Spain. Association for Computational Linguistics.
- Sebastian Gehrmann, Tosin P. Adewumi, Karmanya Aggarwal, Pawan Sasanka Ammanamanchi, Aremu Anuoluwapo, Antoine Bosselut, Khyathi Raghavi Chandu, Miruna-Adriana Clinciu, Dipanjan Das, Kaustubh D. Dhole, Wanyu Du, Esin Durmus, Ondrej Dusek, Chris Emezue, Varun Gangal, Cristina Garbacea, Tatsunori Hashimoto, Yufang Hou, Yacine Jernite, Harsh Jhamtani, Yangfeng Ji, Shailza Jolly, Dhruv Kumar, Faisal Ladhak, Aman Madaan, Mounica Maddela, Khyati Mahajan, Saad Mahamood, Bodhisattwa Prasad Majumder, Pedro Henrique Martins, Angelina McMillan-Major, Simon Mille, Emiel van Miltenburg, Moin Nadeem, Shashi Narayan, Vitaly Nikolaev, Rubungo Andre Niyongabo, Salomey Osei, Ankur P. Parikh, Laura Perez-Beltrachini, Niranjan Ramesh Rao, Vikas Raunak, Juan Diego Rodriguez, Sashank Santhanam, João Sedoc, Thibault Sellam, Samira Shaikh, Anastasia Shimorina, Marco Antonio Sobrevilla Cabezudo, Hendrik Strobelt, Nishant Subramani, Wei Xu, Diyi Yang, Akhila Yerukola, and Jiawei Zhou. 2021. [The GEM benchmark: Natural language generation, its evaluation and metrics](#). *CoRR*, abs/2102.01672.
- David Grangier and Michael Auli. 2018. [QuickEdit: Editing text & translations by crossing words out](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 272–282, New Orleans, Louisiana. Association for Computational Linguistics.

- Hamza Harkous, Isabel Groves, and Amir Saffari. 2020. [Have your text and use it too! end-to-end neural data-to-text generation with semantic fidelity](#). *CoRR*, abs/2004.06577.
- Ari Holtzman, Jan Buys, Maxwell Forbes, and Yejin Choi. 2019. [The curious case of neural text degeneration](#). *CoRR*, abs/1904.09751.
- Edward J. Hu, Yelong Shen, Phillip Wallis, Zeyuan Allen-Zhu, Yuanzhi Li, Shean Wang, and Weizhu Chen. 2021. [Lora: Low-rank adaptation of large language models](#). *CoRR*, abs/2106.09685.
- Melvin Johnson, Mike Schuster, Quoc V. Le, Maxim Krikun, Yonghui Wu, Zhifeng Chen, Nikhil Thorat, Fernanda B. Viégas, Martin Wattenberg, Greg Corrado, Macduff Hughes, and Jeffrey Dean. 2016. [Google’s multilingual neural machine translation system: Enabling zero-shot translation](#). *CoRR*, abs/1611.04558.
- N. Keskar, B. McCann, L. R. Varshney, Caiming Xiong, and R. Socher. 2019. [Ctrl: A conditional transformer language model for controllable generation](#). *ArXiv*, 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](#). *CoRR*, abs/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.
- J. Peter Kincaid, Robert P. Fishburne Jr., Richard L. Rogers, and Brad S. Chissom. 1975. Derivation of new readability formulas (automated readability index, fog count and flesch reading ease formula) for navy enlisted personnel.
- Alon Lavie and Abhaya Agarwal. 2007. [METEOR: An automatic metric for MT evaluation with high levels of correlation with human judgments](#). In *Proceedings of the Second Workshop on Statistical Machine Translation*, pages 228–231, Prague, Czech Republic. Association for Computational Linguistics.
- Brian Lester, Rami Al-Rfou, and Noah Constant. 2021. [The power of scale for parameter-efficient prompt tuning](#). *CoRR*, abs/2104.08691.
- 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.
- Xiang Lisa Li and Percy Liang. 2021. [Prefix-tuning: Optimizing continuous prompts for generation](#).
- 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.
- Pengfei Liu, Weizhe Yuan, Jinlan Fu, Zhengbao Jiang, Hiroaki Hayashi, and Graham Neubig. 2021a. [Pre-train, prompt, and predict: A systematic survey of prompting methods in natural language processing](#).
- Xiao Liu, Yanan Zheng, Zhengxiao Du, Ming Ding, Yujie Qian, Zhilin Yang, and Jie Tang. 2021b. [GPT understands, too](#). *CoRR*, abs/2103.10385.
- Lajanugen Logeswaran, Ann Lee, Myle Ott, Honglak Lee, Marc’Aurelio Ranzato, and Arthur Szlam. 2020. [Few-shot sequence learning with transformers](#). *CoRR*, abs/2012.09543.
- Ilya Loshchilov and Frank Hutter. 2017. [Fixing weight decay regularization in adam](#). *CoRR*, abs/1711.05101.
- L. V. D. Maaten and Geoffrey E. Hinton. 2008. Visualizing data using t-sne. *Journal of Machine Learning Research*, 9:2579–2605.
- Louis Martin, Angela Fan, Éric de la Clergerie, Antoine Bordes, and Benoît Sagot. 2020. [Multilingual unsupervised sentence simplification](#). *CoRR*, abs/2005.00352.
- Amit Moryossef, Ido Dagan, and Yoav Goldberg. 2019. [Improving quality and efficiency in plan-based neural data-to-text generation](#).
- Shashi Narayan, Shay B. Cohen, and Mirella Lapata. 2018. [Don’t give me the details, just the summary! topic-aware convolutional neural networks for extreme summarization](#). *CoRR*, abs/1808.08745.
- Anh Nguyen, Jason Yosinski, Yoshua Bengio, Alexey Dosovitskiy, and Jeff Clune. 2016. [Plug & play generative networks: Conditional iterative generation of images in latent space](#). *CoRR*, abs/1612.00005.
- 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 on Association for Computational Linguistics*, ACL ’02, page 311–318, USA. Association for Computational Linguistics.
- Nivranshu Pasricha, Mihael Arcan, and Paul Buitelaar. 2020. [NUIG-DSI at the WebNLG+ challenge: Leveraging transfer learning for RDF-to-text generation](#). In *Proceedings of the 3rd International Workshop on Natural Language Generation from the Semantic Web (WebNLG+)*, pages 137–143, Dublin, Ireland (Virtual). Association for Computational Linguistics.

- Jeffrey Pennington, Richard Socher, and Christopher Manning. 2014. [GloVe: Global vectors for word representation](#). In *Proceedings of the 2014 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 1532–1543, Doha, Qatar. Association for Computational Linguistics.
- Matthew E. Peters, Sebastian Ruder, and Noah A. Smith. 2019. [To tune or not to tune? adapting pre-trained representations to diverse tasks](#). In *Proceedings of the 4th Workshop on Representation Learning for NLP (RepL4NLP-2019)*, pages 7–14, Florence, Italy. Association for Computational Linguistics.
- Dragomir R. Radev, Rui Zhang, Amrit Rau, Abhinand Sivaprasad, Chiachun Hsieh, Nazneen Fatema Rajani, Xiangru Tang, Aadit Vyas, Neha Verma, Pranav Krishna, Yangxiaokang Liu, Nadia Irwanto, Jessica Pan, Faiaz Rahman, Ahmad Zaidi, Murori Mutuma, Yasin Tarabar, Ankit Gupta, Tao Yu, Yi Chern Tan, Xi Victoria Lin, Caiming Xiong, and Richard Socher. 2020. [DART: open-domain structured data record to text generation](#). *CoRR*, abs/2007.02871.
- Evani Radiya-Dixit and Xin Wang. 2020. [How fine can fine-tuning be? learning efficient language models](#). In *Proceedings of the Twenty Third International Conference on Artificial Intelligence and Statistics*, volume 108 of *Proceedings of Machine Learning Research*, pages 2435–2443, Online. PMLR.
- Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena, Yanqi Zhou, Wei Li, and Peter J. Liu. 2020. [Exploring the limits of transfer learning with a unified text-to-text transformer](#). *Journal of Machine Learning Research*, 21(140):1–67.
- Leonardo F. R. Ribeiro, Martin Schmitt, Hinrich Schütze, and Iryna Gurevych. 2020. Investigating pretrained language models for graph-to-text generation. *arXiv*.
- Thomas Scialom, Louis Martin, Jacopo Staiano, Éric Villemonte de la Clergerie, and Benoît Sagot. 2021. [Rethinking automatic evaluation in sentence simplification](#). *CoRR*, abs/2104.07560.
- Noam Shazeer and Mitchell Stern. 2018. [Adafactor: Adaptive learning rates with sublinear memory cost](#). *CoRR*, abs/1804.04235.
- Matthew Snover, Bonnie Dorr, Richard Schwartz, Linnea Micciulla, and Ralph Weischedel. 2006. A study of translation error rate with targeted human annotation. In *In Proceedings of the Association for Machine Translation in the Americas (AMTA 2006)*.
- Swabha Swayamdipta, Sam Thomson, Kenton Lee, Luke Zettlemoyer, Chris Dyer, and Noah A. Smith. 2018. [Syntactic scaffolds for semantic structures](#). In *EMNLP*.
- Maria Tsimpoukelli, Jacob Menick, Serkan Cabi, S. M. Ali Eslami, Oriol Vinyals, and Felix Hill. 2021. [Multimodal few-shot learning with frozen language models](#). *CoRR*, abs/2106.13884.
- Chris van der Lee, Albert Gatt, Emiel Miltenburg, and Emiel Kraemer. 2021. [Human evaluation of automatically generated text: Current trends and best practice guidelines](#). *Computer Speech & Language*, 67:101151.
- Thomas Wolf, Lysandre Debut, Victor Sanh, Julien Chaumond, Clement Delangue, Anthony Moi, Pierric Cistac, Tim Rault, Rémi 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 M. 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.
- Wei Xu, Courtney Napoles, Ellie Pavlick, Quanze Chen, and Chris Callison-Burch. 2016. Optimizing statistical machine translation for text simplification. *Transactions of the Association for Computational Linguistics*, 4:401–415.
- Dian Yu, Kenji Sagae, and Zhou Yu. 2021. [Attribute alignment: Controlling text generation from pre-trained language models](#). *CoRR*, abs/2103.11070.
- Jingqing Zhang, Yao Zhao, Mohammad Saleh, and Peter J. Liu. 2019. [PEGASUS: pre-training with extracted gap-sentences for abstractive summarization](#). *CoRR*, abs/1912.08777.
- Xingxing Zhang and Mirella Lapata. 2017. Sentence simplification with deep reinforcement learning. *arXiv preprint arXiv:1703.10931*.

A Additional Results

Additional results using the official evaluation scripts for the data-to-text datasets are reported in Tables 5,6,7 to supplement the results in Table 1.

B WebNLG+ 2020 Results

As NLG is notoriously challenging to evaluate, this work assesses model performance on five of the eleven datasets comprising GEM (Gehrmann et al., 2021), a benchmark that intends to provide robust datasets and reproducible standards across an array of NLG tasks. The GEM datasets used in this study are DART, E2E Clean, ASSET, TurkCorpus and WebNLG+ 2020.

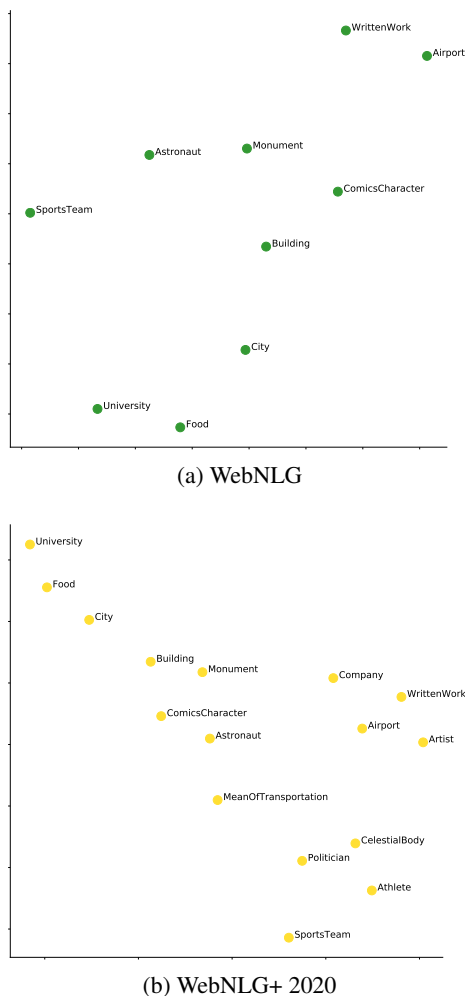


Figure 3: t-SNE visualizations for the encoder constituent of control prefixes representing WebNLG categories seen during training. Each circle represents a category seen during training for the CONTROL PREFIXES (A_1) model. All 15 categories are seen categories in WebNLG+ 2020, along with the category *Company*. WebNLG+ 2020 has 3 additional unseen categories to those shown.

WebNLG+ 2020 is not a component of DART—it was used for the second official WebNLG competition (Castro Ferreira et al., 2020). There are 16 training categories (the 15 categories from WebNLG, but with new examples), alongside 3 unseen categories. Table 8 displays WebNLG+ 2020 results using the same model architectures as used for WebNLG. A similar pattern is revealed, in that CONTROL PREFIXES outperforms prefix-tuning with CONTROL PREFIXES (A_1, A_2) as the top-performing model. This illustrates again the benefit of using both controllable attributes.

In the WebNLG and WebNLG+ 2020 training sets, for the same tripleset, multiple distinct lexicalizations exist. In our experiments, the examples sharing identical tripleset inputs have the same triple order after linearization. This is to aid in comparison with current systems for WebNLG, DART and E2E Clean. Future work would have to assess if architecture-independent improvement in test-set performance can arise by random permutation of the order of triples for training set examples with identical tripleset inputs. The motivation being that this may improve the generalizability of the model, since the model would not learn the order of particular tripleset inputs.

C Prefix-tuning

We make two previously unremarked upon observations of the benefits conferred by using the key-value pair prefix-tuning described in §3.3 compared to prefix-tuning involving augmenting the activations directly (Hu et al., 2021) or prompt-embedding tuning of prompt length ρ . i) The form discussed does not restrict the input length of the base LM. ii) The time complexity at inference time is reduced; for example, if we take a multi-head self-attention computation ($M = N$), the time complexity at inference time is $\mathcal{O}((N + \rho)Nd + Nd^2)$ rather than the greater $\mathcal{O}((N + \rho)^2d + (N + \rho)d^2)$.

D Additional Training Details

All implementations in this study are built on top of the Transformers library (Wolf et al., 2020). As T5 has relative position biases, we set these in all layers pertaining to offsets where the key is part of a prefix to zero. For BART_{LARGE} we adapt the original implementation (Li and Liang, 2021). Table 10 displays the hyperparameters used when training the models reported in this paper.

The general prompt length and each control prompt length are architecture-specific parameters that we choose based on performance on the validation set. We use gradient accumulation across batches to maintain an effective batch size above 64, a linear learning rate scheduler for all models and beam-search decoding. AdamW (Loshchilov and Hutter, 2017) and AdaFactor (Shazeer and Stern, 2018) were used for optimization. We chose the checkpoint with the highest validation score using BLEU for data-to-text, SARI for simplification and ROUGE-2 for summarization. For all tasks, we train our models on single Tesla V100-SXM2-16GB machines, with mixed precision for BART_{LARGE} based models (fp16) and full precision for T5-large based models (fp32).

The CONTROL PREFIXES models with the DART *sub-dataset source* attribute (A_2) use DART as additional data and were trained in two stages: i) on DART, ii) solely on the downstream dataset. The WebNLG prefix-tuning model with DART data shown in Table 10 uses only the human annotated portion of DART. The prefix-tuning models using all of the DART data for WebNLG and E2E Clean were similarly trained in two stages, with identical hyperparameters to CONTROL PREFIXES models using A_2 . Training prefix-tuning on all of DART for WebNLG yielded lower performance than with only the human-annotated DART portion as additional data, so was not reported in Table 1.

Decoding specific parameters were not tuned—we instead mirrored what the top-performing fine-tuned based system used for the particular LM and dataset. For example, a beam width of 5 as in Ribeiro et al. (2020) for T5-large on all data-to-text datasets.

For XSum the source articles are truncated to 512 BPE tokens.

E Simplification Length Control

Fig. 4 depicts the length compression ratio output distribution on the validation set for CONTROL PREFIXES, where a length control prefix of a specific attribute value (0.25,0.5,0.75,1.0) is specified. This clearly demonstrates CONTROL PREFIXES is capable of controlling the target length with respect to the input. Table 11 displays example output generations with each of the 0.25,0.5,0.75,1.0 values specified.

Fig. 5 is supplementary to §6.1, showing all constituents of the length compression control prefixes




Figure 4: Histogram illustrating the influence of different target length ratios on the actual length compression ratio output distribution for the simplification CONTROL PREFIXES model on the TurkCorpus validation set.

for all attribute values. In the WikiLarge training data, there are far fewer training samples where the simplified output is much longer than the complex, original input in WikiLarge. This explains why the representations are not as interpretable for values greater than 1.2.

E.1 QuestEval


The *Gold Reference* results for QuestEval¹⁸ are higher for TurkCorpus compared to ASSET in Table 2. We argue this is because the test set gold references are on average 114 characters for TurkCorpus, as opposed to 98 for ASSET. Therefore, the ASSET references contain less information to answer the generated queries during QuestEval evaluation; and thus, there is lower performance. We argue this shows a limitation with using QuestEval as a reference-less metric for simplification—by favouring longer generations.

F Prefix-tuning + Control Tokens


We propose another architecture ‘prefix-tuning + control tokens’, where all of the original LM parameters, ϕ , still remain fixed, including the embedding matrix. Control has to be exerted through the few control embeddings and prefix-tuning’s ability to steer the frozen ϕ parameters through $< 2\%$ additional parameters. We use this method to inform the model of the same discrete guidance information as in CONTROL PREFIXES, but with control tokens instead of control prefixes.¹⁹ This alter-

¹⁸Although QuestEval can take references, the authors maintain that any improvement in correlation with human performance is very minor.


¹⁹Only the embeddings pertaining to the controllable attributes and the prefix are trained.



(a) Decoder Masked-attention (D_m)



(b) Encoder (E)



(c) Decoder Cross-attention (D_c)


Figure 5: t-SNE visualizations for constituents of the length compression control prefixes learnt as part of the simplification CONTROL PREFIXES model. Each diagram depicts representations of control prefixes corresponding to each length value (41 bins of fixed width 0.05, from 0 to 2) for a particular attention mechanism. The dimension represented on the x-axis is stretched from a 1:1 to 2:1 aspect ratio for labelling clarity.

native method is less expressive than CONTROL PREFIXES, in much the same way as prefix-tuning is more expressive than prompt-embedding tuning. Prefix-tuning + control tokens also does not benefit from the shared re-parameterizations (§3.3) that we argue allow for more effective demarcation of control of the fixed LM in each attention class subspace.


Table 9 reveals that CONTROL PREFIXES outperforms prefix-tuning + control tokens on the data-to-text and summarization datasets, while the results are both comparable to the *Gold References* on simplification datasets. This indicates that CONTROL PREFIXES is better able to integrate and leverage

guidance signal at the input-level, whilst maintaining the *fixed-LM* property, than prefix-tuning + control tokens.

G Varying Prompt Length



(a) BART_{LARGE}



(b) T5-large

Figure 6: Prefix-tuning results of a model parameter search on several datasets for the optimal prompt length per dataset. These results are for the metric monitored per task on the respective validation sets indicated in the legend. $\phi\%$ denotes the % of additional parameters to the number of fixed-LM parameters required at inference time. The y -axis is a relative measure: the validation set performance as a % of the maximum attained in the parameter search.

Our research is not solely focused on parameter efficiency, but also on the effectiveness of adapting an already parameter efficient, fixed-LM method (adding <3% additional parameters). The only way to add parameters with prefix-tuning is to increase the prompt length. XSum is the only dataset considered where performance does not plateau when increasing prompt length²⁰, therefore we ensure

²⁰We do not observe performance degradation, such as

CONTROL PREFIXES does not have more parameters than prefix-tuning to ensure a fair comparison.

The only way to add parameters with prefix-tuning is by increasing prompt length. Fig. 6 illustrates how performance saturation is observed—after a certain prompt length performance plateaus. Different datasets require varying prompt lengths to attain near maximum performance in a parameter search for prompt length. For the data-to-text datasets, near maximum performance (>99% of the maximum validation score in the search) is reached with a prompt length of 1 or 2.

H Qualitative Examples

For data-to-text, Table 13 displays example CONTROL PREFIXES output for WebNLG input belonging to unseen categories, along with the zero-shot procedure. Table 13 depicts example CONTROL PREFIXES (A_1, A_2) output alongside prefix-tuning model output for WebNLG+ 2020 input. For simplification, Table 12 compares the fixed-LM guided generations of CONTROL PREFIXES to the fine-tuned BART_{LARGE} with ACCESS (Martin et al., 2020). For summarization, Table 15 depicts cherry-picked CONTROL PREFIXES generated summaries for XSum input, alongside T5-large fine-tuned summaries that have higher ROUGE scores. This is to illustrate how CONTROL PREFIXES can achieve higher human assessment through GENIE than top-performing fine-tuned models, whilst attaining lower automatic metric scores.

described by Hu et al. (2021), when utilizing different forms of prefix-tuning. This is shown in G.

	$\phi\%$	BLEU	METEOR	TER ↓	BERTScore(F1)
T5-large fine-tuned*	100	50.66	40	43	0.95
Prefix-tuning	1.0	51.20	40.62	43.13	0.95
CONTROL PREFIXES (A_1)	1.1	51.95	41.07	42.75	0.95

Table 5: Detailed results on the DART test set to complement Table 1. T5-large fine-tuned is the current SOTA (Radev et al., 2020). We report results on the official evaluation script for v1.1.1, the same version as the official leaderboard, available here: <https://github.com/Yale-LILY/dart>. *Results for this model were only reported to the significant figures shown. $\phi\%$ denotes the % of additional parameters to the number of fixed-LM parameters required at inference time.

	$\phi\%$	BLEU			METEOR			TER ↓		
		S	U	A	S	U	A	S	U	A
T5-large	100	64.89	54.01	59.95	46	43	44	34	41	37
SOTA	100	65.82	56.01	61.44	46	43	45	32	38	35
Prefix-tuning	1.0	66.95	55.39	61.73	46.73	42.71	44.87	31.34	39.01	34.86
CONTROL PREFIXES (A_1)	1.4	67.32	55.38	61.94	46.78	42.77	44.92	30.96	39.01	34.65
+Data: DART										
Prefix-tuning	1.0	67.05	55.37	61.78	46.69	42.82	44.90	31.36	38.79	34.77
CONTROL PREFIXES (A_2)	1.0	66.99	55.56	61.83	46.67	42.87	44.91	31.37	38.53	34.65
CONTROL PREFIXES (A_1, A_2)	1.4	67.15	56.41	62.27	46.64	43.18	45.03	31.08	38.78	34.61

Table 6: Detailed results on the WebNLG test set to complement Table 1. S, U and A refer to the *Seen*, *Unseen* and *All* portions of the WebNLG dataset. Several of the baseline results were only reported to the significant figures shown.

	$\phi\%$	BLEU	NIST	METEOR	R-L	CIDEr
T5-Large	100	41.83	6.41	0.381	56.0	1.97
SOTA	100	43.6	-	0.39	57.5	2.0
Prefix-tuning	1.0	43.66	6.51	0.390	57.2	2.04
+Data: DART						
Prefix-tuning	1.0	43.04	6.46	0.387	56.8	1.99
CONTROL PREFIXES (A_2)	1.0	44.15	6.51	0.392	57.3	2.04

Table 7: Detailed results on the E2E Clean test set to complement Table 1. The SOTA baseline result was only reported to the significant figures shown.

	$\phi\%$	BLEU	METEOR	chrF++	TER ↓	BLEURT
T5-large* [†]	100	51.74	0.403	0.669	0.417	0.61
Prefix-tuning	1.0	54.74	0.417	0.693	0.399	0.62
CONTROL PREFIXES (A_1)	1.6	54.97	0.417	0.693	0.398	0.62
+Data: DART						
CONTROL PREFIXES (A_2)	1.0	54.92	0.418	0.695	0.397	0.62
CONTROL PREFIXES (A_1, A_2)	1.6	55.41	0.419	0.698	0.392	0.63

Table 8: **WebNLG+ 2020**. The overall WebNLG+ 2020 test set results using the official evaluation script. *As the model outputs are publicly available, we are able to run evaluation to achieve the same precision. [†]Results from Pasricha et al. (2020), who before fine-tuning on the WebNLG+ data, further pre-train T5-large using a Mask Language Modelling objective (with 15% of the tokens masked) on the WebNLG corpus and a corpus of DBpedia. A_1 signifies models trained with control prefixes for the *WebNLG category* attribute, and A_2 with control prefixes for the DART *sub-dataset source* attribute.

	DART	WebNLG	E2E Clean	ASSET		TurkCorpus		XSum
	BLEU			SARI	QuestEval	SARI	QuestEval	R-2
Prefix-tuning + Control Tokens	51.72	61.89	43.57	43.64	0.63	42.36	0.66	20.70
CONTROL PREFIXES	51.95	62.27	44.15	43.58	0.64	42.32	0.66	20.84

Table 9: **Prefix-tuning + Control Tokens.** Comparison of our best CONTROL PREFIXES model for each dataset with prefix-tuning + control tokens for the same attributes. The guided simplification models are the average test set results over 5 random seeds.

Model	Stage	L-rate	Opt	Warmup-steps	Epochs	Batch Size	Effective Batch	Beam Width	LN- α	Min Target	Max Target	No Repeat Trigram
<i>DART (T5-large)</i>												
Prefix-tuning	-	7e-5	Ada	2000	40	6	96	5	1	0	384	No
CONTROL PREFIXES (A_1)	-	7e-5	Ada	2000	40	6	96	5	1	0	384	No
<i>E2E Clean (T5-large)</i>												
Prefix-tuning	-	8e-5	Ada	2000	50	6	96	5	1	0	384	No
CONTROL PREFIXES (A_2)	1	7e-5	Ada	2000	30	6	96	5	1	0	384	No
	2	5e-5	Ada	2000	50	6	96	5	1	0	384	No
<i>WebNLG (T5-large)</i>												
Prefix-tuning	-	7e-5	Ada	2000	30	6	96	5	1	0	384	No
CONTROL PREFIXES (A_1)	-	7e-5	Ada	2000	40	6	96	5	1	0	384	No
+Data: DART												
Prefix-tuning	-	7e-5	Ada	2000	40	6	96	5	1	0	384	No
CONTROL PREFIXES (A_2)	1	7e-5	Ada	2000	30	6	96	5	1	0	384	No
	2	3e-5	Ada	2000	30	6	96	5	1	0	384	No
CONTROL PREFIXES (A_1, A_2)	1	7e-5	Ada	2000	30	6	96	5	1	0	384	No
	2	3e-5	Ada	2000	30	6	96	5	1	0	384	No
<i>XSum (BART_{LARGE})</i>												
Prefix-tuning	-	7e-5	AdamW	2000	40	8	128	6	1	10	60	✓
CONTROL PREFIXES (A_1, A_2)	-	7e-5	AdamW	2000	40	8	128	6	1	10	60	✓
<i>ASSET & TurkCorpus (BART_{LARGE})</i>												
Prefix-tuning	-	5e-5	AdamW	2000	30	8	64	6	0.8	3	100	✓
CONTROL PREFIXES	-	4e-5	Ada	5000	30	8	64	6	1	3	100	✓

Table 10: **Hyperparameters.** Detailed hyperparameter reporting for the models in this work. If the training procedure is multi-stage, each stage is indicated. L-rate is the learning rate, all learning follows a linear learning rate scheduler; Opt refers to the optimizer, Ada (Adafactor) or AdamW; Effective Batch = Batch size x # of gradient accumulation batches; LN- α refers to the α in length normalization during beam search.

ASSET Corpus	
	Source: The West Coast blues is a type of blues music characterized by jazz and jump blues influences, strong piano-dominated sounds and jazzy guitar solos, which originated from Texas blues players relocated to California in the 1940s.
Gold Reference [†]	The West Coast blues has jazz and jump blues influences. It also has piano-dominated sounds and jazzy guitar solos, which originated from Texas blues players who moved to California in the 1940s.
CONTROL PREFIXES ($A_{1,4}$) Length Control: 0.25	West Coast blues is a type of blues music from the United States.
CONTROL PREFIXES ($A_{1,4}$) Length Control: 0.5	The West Coast blues is a type of blues music made from Texas blues players who moved to California in the 1940s. It has jazz and jump blues influences.
CONTROL PREFIXES ($A_{1,4}$) Length Control: 0.75	The West Coast blues is a type of blues music that originated in Texas in the 1940s. It is characterized by jazz and jump blues influences, strong piano-dominated sounds and jazzy guitar solos.
CONTROL PREFIXES ($A_{1,4}$) Length Control: 1.0	The West Coast blues is a type of blues music characterized by jazz and jump blues influences, strong piano-dominated sounds and jazzy guitar solos, which originated from Texas blues players who moved to California in the 1940s.

Table 11: **Controlled text simplification:** a qualitative example from the ASSET Corpus validation set, with our corresponding CONTROL PREFIXES model outputs. Only the control prefix for the length control ratio is varied, depicted in red. [†]Note, this is one random gold reference out of 10.

ASSET Corpus	
	Source: The Great Dark Spot is thought to represent a hole in the methane cloud deck of Neptune.
Gold Reference [†]	The Great Dark Spot represents a hole in the methane cloud of Neptune.
CONTROL PREFIXES	It is thought that the Great Dark Spot is a hole in Neptune’s methane cloud deck.
BART _{LARGE} with ACCESS	The Great Dark Spot looks like a hole in the methane cloud deck of Neptune.
	Source: Fives is a British sport believed to derive from the same origins as many racquet sports.
Gold Reference [†]	Fives is a British sport developed from the same origins as many racquet sports.
CONTROL PREFIXES	Fives is a British sport. It is believed to have its origins in racquet sports.
BART _{LARGE} with ACCESS	Fives is a British sport. It is thought to come from the same as many racquet sports.
	Source: Nevertheless, Tagore emulated numerous styles, including craftwork from northern New Ireland, Haida carvings from the west coast of Canada (British Columbia), and woodcuts by Max Pechstein.
Gold Reference [†]	Tagore copied many styles. These included craftwork from northern New Ireland, Haida carvings from western Canada and woodcuts by Max Pechstein.
CONTROL PREFIXES	Tagore emulated many different styles of art, including Haida carvings from the west coast of Canada (British Columbia), and woodcuts by Max Pechstein.
BART _{LARGE} with ACCESS	Tagore copied many styles. He copied craftwork from northern New Ireland, Haida carvings from the west coast of Canada (British Columbia), and woodcuts by Max Pechstein.

Table 12: **Fixed-LM vs fine-tuned controlled text simplification.** CONTROL PREFIXES and BART_{LARGE} with ACCESS (Martin et al., 2020) generated simplifications chosen from the ASSET Corpus test set. [†]Note, this is one random gold reference out of 10 for each example. The examples shown for CONTROL PREFIXES and BART_{LARGE} with ACCESS are also randomly selected from one of the five model outputs.

WebNLG	
Unseen Category: <i>Athlete</i> Zero-shot -> <i>SportsTeam</i>	Source: <H> FC Torpedo Moscow <R> season <T> 2014-15 Russian Premier League <H> Aleksandr Chumakov <R> club <T> FC Torpedo Moscow <H> FC Torpedo Moscow <R> manager <T> Valery Petrakov <H> FC Torpedo Moscow <R> chairman <T> Aleksandr Tukmanov
Gold	Valery Petrakov is the manager of FC Torpedo Moscow and its chairman is Aleksandr Tukmanov. Aleksandr Chumakov plays for the club which spent the 2014-15 season in the Russian Premier League.
Prefix-tuning	Aleksandr Tukmanov and Valery Petrakov are the managers of FC Torpedo Moscow. The club played in the Russian Premier League in 2014-15 and their chairman is Aleksandr Tukmanov.
CONTROL PREFIXES (A_1)	Aleksandr Chumakov plays for FC Torpedo Moscow which is managed by Valery Petrakov. The club's chairman is Aleksandr Tukmanov and they played in the Russian Premier League in the 2014-15 season.
<hr/>	
Unseen Category: <i>MeanOfTransportation</i> Zero-shot -> <i>Airport</i>	Source: <H> Costa Crociere <R> location <T> Genoa <H> Costa Crociere <R> parent Company <T> Carnival Corporation & plc <H> AIDAstella <R> operator <T> AIDA Cruises <H> AIDAstella <R> builder <T> Meyer Werft <H> AIDAstella <R> owner <T> Costa Crociere
Gold	Carnival Corporation & plc is the parent company of Costa Crociere in Genoa, who own the AIDAstella. AIDAstella was built by Meyer Werft and is operated by AIDA Cruises.
Prefix-tuning	Costa Crociere is located in Genoa and is owned by Carnival Corporation & plc. AIDAstella is operated by AIDA Cruises and was built by Meyer Werft.
CONTROL PREFIXES (A_1)	Costa Crociere is located in Genoa and is owned by AIDA Cruises. AIDAstella was built by Meyer Werft and is operated by AIDA Cruises. The parent company of Costa Crociere is Carnival Corporation & plc.

Table 13: **WebNLG example generations:** sources are shown in their linearized form, as fed to the T5-large based models, with prefix-tuning output and one of the gold references shown for comparison with CONTROL PREFIXES output. Triplesets are from WebNLG unseen categories and the zero-shot procedure is depicted using the textual category labels. As an example, for the unseen category *Athlete*, the closest Glove embedding belonging to a *seen* category label in embedding space is *SportsTeam*. Therefore the trained control prefix relating to *SportsTeam* is used for this example at inference time.

WebNLG+ 2020	
WebNLG MeanOfTransportation (Seen with Unseen Entities)	Source: <H> Pontiac Rageous <R> production Start Year <T> 1997 <H> Pontiac Rageous <R> assembly <T> Michigan <H> Pontiac Rageous <R> assembly <T> Detroit <H> Pontiac Rageous <R> production End Year <T> 1997 <H> Pontiac Rageous <R> body Style <T> Coupe <H> Pontiac Rageous <R> manufacturer <T> Pontiac
Gold	The Pontiac Rageous was a car with a coupe body style manufactured by Pontiac. Assembled in both Michigan and Detroit, it went into production in 1997, ending in the same year.
Prefix-tuning	The Pontiac Rageous is a coupe manufactured by Pontiac. It is assembled in Detroit, Michigan and began production in 1997.
CONTROL PREFIXES (A_1, A_2)	The Pontiac Rageous is manufactured by Pontiac in Detroit, Michigan. Its production began in 1997 and ended in 1997. The Pontiac Rageous has a coupe body style.
<hr/>	
WebNLG (Unseen) Unseen Category: <i>MusicalWork</i> Zero-shot -> <i>Artist</i>	Source: <H> Bootleg Series Volume 1: The Quine Tapes <R> genre <T> Rock music <H> Bootleg Series Volume 1: The Quine Tapes <R> preceded By <T> Squeeze The Velvet Underground album <H> Bootleg Series Volume 1: The Quine Tapes <R> record Label <T> Polydor Records <H> Bootleg Series Volume 1: The Quine Tapes <R> recorded In <T> San Francisco
Gold	The Velvet Underground Squeeze album was succeeded by the rock album Bootleg Series Volume 1: The Quine Tapes, recorded under record label Polydor Records in San Francisco.
Prefix-tuning	The record label of Bootleg Series Volume 1: The Quine Tapes is Polydor Records. It was recorded in San Francisco and was preceded by Squeeze The Velvet Underground. Its genre is rock music.
CONTROL PREFIXES (A_1, A_2)	Squeeze The Velvet Underground was preceded by Bootleg Series Volume 1: The Quine Tapes, which was recorded in San Francisco and released by Polydor Records. The genre of the album is rock music.

Table 14: **WebNLG+ 2020 generations:** sources are shown in their linearized form as fed to the T5-large based models. The DART sub-dataset *Source* control prefix is highlighted, along with the final *Category* control prefix. The zero-shot procedure is depicted for the Unseen Category *MusicalWork*. The closest embedding belonging to a *Seen* category in embedding space is *Artist*.

XSum	
news world	Kamal C Chavara was detained by the police in Kerala state on Sunday after the youth wing of the Hindu nationalist BJP lodged a complaint against him. Last month, the Supreme Court ruled that the anthem must be played in every cinema before a film is screened. Some 20 people have been held in Kerala and Tamil Nadu since then for remaining seated during the anthem. Also, India's colonial-era sedition law has been often used against students, journalists, writers and social activists and those critical of the government. Reports said that the BJP's youth wing lodged a complaint against a Facebook post by Mr Chavara which allegedly insulted the anthem. The post was apparently an excerpt from one of his books. Senior police official Sateesh Bino told the NDTV news channel that the writer-activist "is being questioned for his controversial post on the national anthem on Facebook" and had been charged with sedition. Earlier this month, 12 people were arrested at a cinema in Kerala, after they remained seated while the national anthem played. The cinemagoers, who were attending an international film festival, were later freed but they face charges of "failure to obey an order issued by a public servant, thereby causing obstruction or annoyance to others". And at a cinema in Chennai, eight people who did not stand for the anthem were assaulted and abused, police said. The eight were later charged with showing disrespect to the anthem.
Gold	A writer in India has been charged with sedition for allegedly showing disrespect to the national anthem.
T5-large fine-tuned (70.97/48.28/70.97)	A prominent Indian writer has been charged with sedition for defying the National Anthem.
CONTROL PREFIXES (A_1, A_2) (59.46/34.29/54.05)	An Indian writer-activist has been charged with sedition over a post on Facebook which allegedly insulted the national anthem.
sport horse-racing	The 33-1 shot, ridden by David Mullins and trained by Mouse Morris, triumphed at Aintree in April to become the first novice to win the race since 1958. The nine-year-old, owned by the Gigginstown House Stud, has twice recovered from a cracked pelvis. "We didn't want to send him back to Aintree with a big weight, that wouldn't be fair," said Gigginstown's racing manager Eddie O'Leary. "He provided us with our first Grand National and we'll never forget him." BBC horse racing correspondent Cornelius Lysaght: "As the first Grand National winner for owner Michael O'Leary's burgeoning Gigginstown House Stud as well as the first novice chaser to win the race in nearly 60 years, Rule The World has his place in history. "Though he ran highly respectably at Punchestown after Aintree, O'Leary had already hinted that, having defied serious injury to reach one of the great pinnacles, he had perhaps done his bit. "What a season for Gigginstown, with success at Aintree, in the Irish National and Cheltenham Gold Cup, but at a price. Rule the World has been retired and there are doubts whether Gold Cup winner Don Cossack will race again."
Gold	This year's Grand National winner Rule The World has been retired.
T5-large fine-tuned (57.14/46.15/57.14)	A Grand National-winning novice ridden by the brilliant rider Rule The World has been retired.
CONTROL PREFIXES (A_1, A_2) (55.17/44.44/55.17)	Winning Grand National hurdler Rule the World has been retired from racing at the age of nine.

Table 15: **XSum generated summaries** for T5-large fine-tuned and CONTROL PREFIXES based on BART_{LARGE}. These are presented alongside the source document and the sole gold reference. Source documents are truncated to 300 words if necessary. **ROUGE-1/ROUGE-2/ROUGE-L** are reported in bold. The **news/sport** control prefix and the related **sub-directory** control prefix are highlighted.