

# Extracting Latent Steering Vectors from Pretrained Language Models

Nishant Subramani<sup>†</sup> Nivedita Suresh<sup>◇</sup> Matthew E. Peters<sup>†</sup>

<sup>†</sup>Allen Institute for Artificial Intelligence, Seattle, WA, USA

<sup>◇</sup>Arrive Bio, San Francisco, CA, USA

{nishants, matthewp}@allenai.org

{nive}@arrivebio.com

## Abstract

Prior work on controllable text generation has focused on *learning* how to control language models through trainable decoding, smart-prompt design, or fine-tuning based on a desired objective. We hypothesize that the information needed to steer the model to generate a target sentence is already encoded within the model. Accordingly, we explore a different approach altogether: *extracting* latent vectors directly from pretrained language model decoders without fine-tuning. Experiments show that there exist *steering vectors*, which, when added to the hidden states of the language model, generate a target sentence nearly perfectly (> 99 BLEU) for English sentences from a variety of domains. We show that vector arithmetic can be used for unsupervised sentiment transfer on the Yelp sentiment benchmark, with performance comparable to models tailored to this task. We find that distances between steering vectors reflect sentence similarity when evaluated on a textual similarity benchmark (STS-B), outperforming pooled hidden states of models. Finally, we present an analysis of the intrinsic properties of the steering vectors. Taken together, our results suggest that frozen LMs can be effectively controlled through their latent steering space.<sup>1</sup>

## 1 Introduction

Leveraging large pretrained language models trained on massive Web corpora has become the go-to approach to solve natural language processing tasks (Peters et al., 2018; Radford et al., 2018; Devlin et al., 2018; Brown et al., 2020). As a result, controlling these models has become paramount as many applications of NLP technology require control over the generations of the model. Prior work aims to *learn* how to control language models and falls in three categories: trainable decoding (Gu

et al., 2017; Deng et al., 2020), smart-prompt design (Shin et al., 2020; Lester et al., 2021), and fine-tuning based on a desired objective (Krause et al., 2021; Weng, 2021). Further, many works opt to train auto-encoder based models for controllable text generation (Shen et al., 2017, 2020; Mai et al., 2020). These approaches make controllability easier by learning a latent space that is more easily manipulated to encourage models to generate text corresponding to a target attribute such as positive sentiment in the case of sentiment transfer.

We take a more direct approach and explore whether it is possible to *extract* latent vectors directly from pretrained language model decoders without fine-tuning. We call these vectors *steering vectors* and define the *latent steering space* of a sentence under a language model by the set of extracted steering vectors, which steer the model to generate that sentence exactly. During decoding, we add our steering vector to the hidden states of the language model to generate the target sentence. Rather than training a model to learn steering vectors, we provide several methods to *extract* fixed-length steering vectors directly from pretrained language model decoders. Experiments show that we can extract steering vectors effectively, achieving nearly perfect recovery for English sentences from a variety of domains without fine-tuning the underlying language model at all.

Next, we take our extracted steering vectors and explore whether they can be used for unsupervised sentiment transfer on the Yelp sentiment benchmark (Zhang et al., 2015). We find that adding an offset vector to extracted steering vectors performs comparably to carefully designed, autoencoder-based models. To see whether steering vectors encode semantics, we explore whether they can be used for unsupervised textual similarity. On the semantic textual similarity benchmark (STS-B, Cer et al. (2017)), our steering vectors outperform extractive methods such as averaging language model

<sup>1</sup>Code is available at [https://github.com/nishantsubramani/steering\\_vectors](https://github.com/nishantsubramani/steering_vectors).

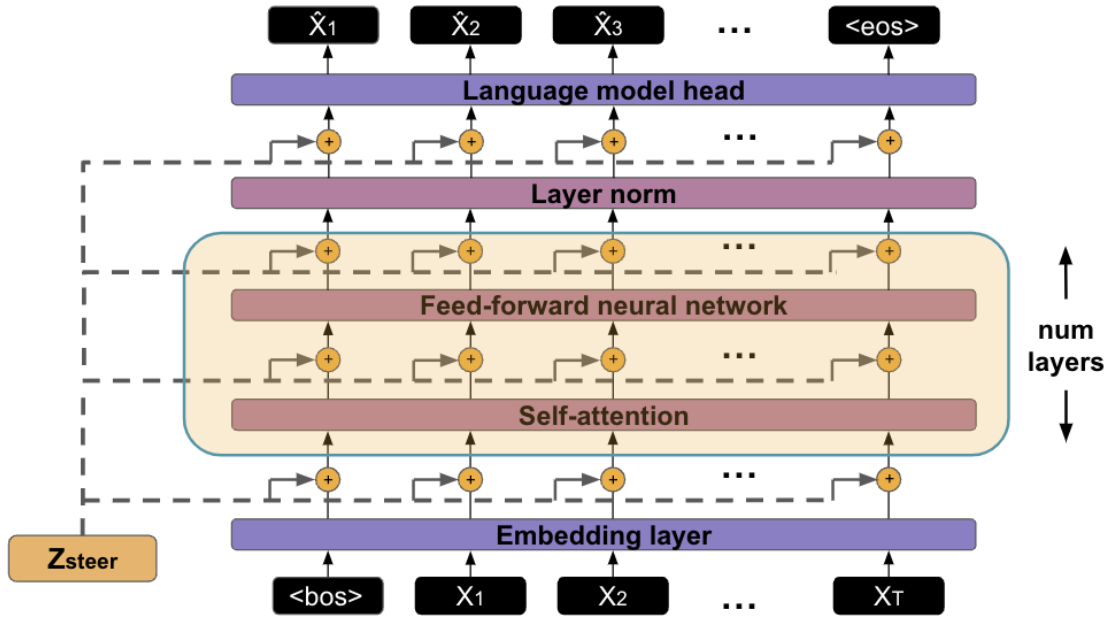


Figure 1: Our approach adds a vector  $z_{steer}$  to the activations of a pretrained transformer decoder to steer it to decode a desired target sentence. We experiment with adding  $z_{steer}$  to different locations inside a GPT-2 model at different timesteps. Experiments reveal that our approach can recover sequences nearly perfectly and that injecting the steering vector in the middle layers of the transformer stack performs best. Layer normalizations and residual connections inside the transformer block are omitted for clarity.

hidden states and GloVe vectors (Pennington et al., 2014) when measuring the cosine similarity between vectors, but fall short of lexical methods tailored to semantic similarity tasks and methods that finetune on natural language inference datasets.

Lastly, we analyze the intrinsic properties of the latent space of our steering vectors. Experiments show that decoding from interpolations in the latent space produces meaningful output, and that steering vectors from different domains cluster together. Also, we find that our methods do not simply memorize the target sequence like a naive compression algorithm, and instead leverage the model. Taken together, our results suggest that frozen language models can be controlled effectively through their latent steering space.

## 2 Extracting Steering Vectors

This section discusses our method for extracting a steering vector for a target sentence from a frozen, pretrained language model. Throughout this paper, we use GPT2 as our language model and use its 117M parameter model size (Radford et al., 2019), although our approach can be directly applied to any transformer-based autoregressive language model decoder (Vaswani et al., 2017).

### 2.1 Steering Vectors

In controllable text generation and textual style transfer, prior work based on denoising and variational autoencoders opt for a disentangling approach. These approaches encode the source sequence into a fixed-length vector using an encoder, apply style transformations using a controller, and finally decode from the transformed vector using a decoder (Shen et al., 2017; Jin et al., 2020). Instead of learning an encoder and controller to uncover a representation, we ask whether its possible to extract a vector directly from a pretrained language model decoder in order to steer the model.

Due to the success of hidden layer manipulations for language models including adapter-based fine-tuning (Houlsby et al., 2019), plug-and-play language models (Dathathri et al., 2019), and offset-vector-based recovery and style transfer among others (Subramani et al., 2019; Shen et al., 2020; Mai et al., 2020; Montero et al., 2021), we choose to manipulate the hidden states as well.

Our method works by adding a fixed-length vector  $z_{steer}$  to the hidden states of a pretrained and frozen LM. For a desired target sentence, we randomly initialize  $z_{steer}$  and optimize it via gradient descent to maximize the likelihood of the model given the target sentence. At decoding time, we

feed a  $z_{steer}$  to the model and perform decoding as usual. The choice of a fixed-length vector makes analysis more meaningful, allowing us to compare vectors for different sentences with different lengths in the same representation space.

## 2.2 Discovering steering vectors

We define our steering vectors  $z_{steer} \in \mathbb{R}^{d'}$ . In our experiments,  $d' \leq d$ , where  $d$  is the hidden dimension of the underlying language model (for GPT2-117M,  $d = 768$ ). If  $d' < d$ , we project  $z_{steer}$  using a semi-orthogonal matrix,  $W_{steer} \in \mathbb{R}^{d' \times d}$ , which preserves scale.  $W_{steer}$  is initialized randomly, never trained, and never updated.

We estimate a steering vector  $\hat{z}_{steer} \in \mathbb{R}^{d'}$  via the language model for a sentence  $x$  by maximizing the log probability of  $x$ , while keeping the language model fixed:

$$\hat{z}_{steer} = \operatorname{argmax}_{z_{steer} \in \mathcal{Z}} \sum_{t=1}^T \log p(x_t | \mathbf{x}_{<t}, z_{steer}) \quad (1)$$

Here,  $\mathcal{Z} \in \mathbb{R}^{d'}$ . Note: we find a single steering vector  $z_{steer}$  for each sentence  $x$ . We use stochastic gradient descent with the Adam (Kingma and Ba, 2014) optimizer and cross entropy loss to find the best  $\hat{z}_{steer}$ , while freezing the language model. See algorithm 1 for the pseudocode.

Since our method adds  $z_{steer}$  to the activations of the model, the layer we add  $z_{steer}$  to affects recoverability. We experiment with injecting  $z_{steer}$  at different layers (*injection locations*): at the embedding layer, right before language model head (LM Head), after self-attention layer in the transformer stack, after feed-forward layer in the transformer stack as well as combinations of them. In addition to varying injection locations, we also vary the timesteps where  $z_{steer}$  gets added. We experiment with adding  $z_{steer}$  at just the first timestep and at every timestep. See Figure 1 for details.

## 2.3 Steering Language Models

We steer the language model using  $z_{steer}$  to generate a target sentence  $x$  by passing in a beginning-of-sentence token and  $z_{steer}$  to the model. Since we are interested in exact generation, all results presented use greedy decoding without assuming a true length. We stop when decoding produces an end-of-sentence token or produces 1024 tokens, the maximum length that GPT-2 can generate.

---

### ALGORITHM 1: Extracting $z_{steer}$ for a sentence

---

**Input** :  $x$  – target sentence  
 $M$  – pretrained language model  
 $\theta$  – pretrained language model weights  
 $I_L$  – injection location  
 $I_T$  – injection timestep  
 $d$  – dimension of  $z_{steer}$

**Output** :  $z_{steer}$  – extracted candidate steering vector

```

1  $z_{steer} \sim \text{xavier\_normal}(d)$ 
2 for  $i \leftarrow [1, 2, \dots, N]$  do
3    $\text{logits} = M_{\theta}.\text{forward}(x, z_{steer}, I_L, I_T)$ 
4    $\mathcal{L} = \text{XENT}(\text{logits}, x)$ 
5    $\mathcal{L}.\text{backward}()$ 
6    $z_{steer} = z_{steer} + lr * \frac{\partial \mathcal{L}}{\partial z_{steer}}$ 
7 end
8 return  $z_{steer}$ 

```

---

## 3 Can we extract steering vectors?

Here, we show that we can robustly extract steering vectors that generate target sentences perfectly.

### 3.1 Experimental setup

We gather a broad corpus spanning four different domains and measure the extent to which our approach can extract a steering vector for each sentence under a variety of experimental conditions, where we vary injection locations and timesteps.

**Data Collection** For these experiments on sentence recoverability, we create a dataset which combines four corpora from different domains: movie dialogs (movies), classic books (books), news articles (news), and Wikipedia (wiki). For movies, we choose the Cornell Movie Dialogs corpus (Danescu-Niculescu-Mizil and Lee, 2011), which consists of fictional conversations from movie scripts. We choose NLTK’s Gutenberg dataset for our books portion, which consists of a subset of texts from Project Gutenberg (Lebert, 2008). Our news subset comes from the Gigaword dataset for abstractive summarization (Graff et al., 2003). Lastly, our Wikipedia portion comes from WikiText-103 (Merity et al., 2017). For movies, news, and wiki, we extract sentences from its pre-specified validation set. For books, since NLTK’s Gutenberg dataset lacks a pre-specified data split, we consider the entire dataset.

**Data Preprocessing** We sentence tokenize all datasets using NLTK’s sentence tokenizer. To construct our dataset, we group sentences by sentence length into 8 bins: 5-10, 10-15, 15-20, 20-25, 25-30, 30-35, 35-40, and 40-128 using NLTK’s word-level, regular expression tokenizer. Next, we ran-

Injection location	Timestep	BLEU-4
Embedding	all timesteps	33.99
<b>Layer 6 (self attn)</b>	<b>all timesteps</b>	<b>100.0</b>
<b>Layer 6 (self attn)</b>	<b>first timestep</b>	<b>99.80</b>
<b>Layer 7 (feed fwd)</b>	<b>all timesteps</b>	<b>100.0</b>
<b>Layer 7 (feed fwd)</b>	<b>first timestep</b>	<b>99.25</b>
<b>All layers (self attn + feed fwd)</b>	<b>all timesteps</b>	<b>100.0</b>
<b>All layers (self attn + feed fwd)</b>	first timestep	91.72
LM head	all timesteps	6.72

Table 1: Sentence recovery for steering vectors when injected into different layers of the transformer model (Figure 1) and at multiple timesteps (all timesteps or first timestep). Results show that injecting a steering vector into the transformer stack, even at just the first timestep, can lead to nearly perfect recovery as long as it is in the middle of the network (layers 6 or 7 of 12).

domly sample 8 sentences from each bin to examine the efficacy of our method for a variety of sequence lengths.

**Measuring the Effectiveness of Steering** Given a target sentence  $s$ , we measure how well the steering vector  $z_{steer}$  can recover the target sentence by first greedily decoding from the language model with  $z_{steer}$ , and then computing smoothed BLEU-4 using the target sentence  $s$  and our decoded reconstruction  $\hat{s}$  (Papineni et al., 2002; Chen and Cherry, 2014).

**Hyperparameter Search** Our initial experiments showed little variation to most hyperparameters such as initialization method and learning rate schedule, so we fixed them in subsequent experiments using the values in Table 6 in the appendix. We choose GPT2-117M as our language model and evaluate recovery on our dataset while varying injection locations and injection timesteps, the two hyperparameters that affect results significantly. We present a subset of the results in Table 1 and the full set in the appendix (Tables 7, 8, and 9).

### 3.2 Recovery effectiveness

Table 1 shows reconstruction performance for several injection methods and indicates that we can recover a target sentence with perfect recovery when injecting  $z_{steer}$  in the middle of the transformer stack (layers 6 or 7 of 12) at just the first timestep and at all timesteps, for sequences up to 128 to-




Figure 2: TSNE projection of 8 steering vectors initialized from different random seeds for 20 different sentences (injected at layer 6, after self-attention).  $z_{steer}$  is well-separated for different sentences, and the different seeds are tightly clustered for the same target sentence, indicating that our extraction method is robust.

kens. We surmise that the middle layers of the transformer stack encode sufficiently rich feature representations that a small perturbation of a hidden layer, a steering vector, is sufficient to recover a sentence. The success of steering vectors when injected in the middle of the transformer could help explain why adapter-based fine-tuning is effective.

In contrast, we find that we cannot steer GPT-2 at either the embedding or final language model head locations. We suspect this is due to the fact that the embedding layer solely captures low-level information (Lin et al., 2019; Ethayarajh, 2019; Rogers et al., 2020). Poor recovery at the LM head location is somewhat surprising, but could be explained by noting that the model has very low capacity above this layer. This suggests that alternative steering mechanisms, such as DExperts, that intervene at the output layers could potentially be improved by modifying hidden states elsewhere in the transformer stack (Liu et al., 2021).

**Robustness** Now that we have established that steering vector extraction is possible, we explore whether there exist multiple steering vectors which recover the same sentence, and if so, what the relationship is between these vectors. To do this, we take all 64 sentences from the books subset of the main dataset and initialize 8 different steering vectors for each sentence from different seeds. Experiments reveal that for most sentences (63 of 64) all initializations recover the target sentence

perfectly, confirming the robustness of our method.

Latent geometry in text-based auto-encoders struggle with mapping vectors from one space to another consistently (e.g. token space to latent space) (Bowman et al., 2016; Shen et al., 2020). The denoising auto-encoder offers a more consistent token space to latent space mapping (Vincent et al., 2008). To explore whether our steering vectors have a distance-preserving mapping, we cluster the different initializations of steering vectors. We extract 8 steering vectors for each of 20 sentences from the books corpus and down-project them into two-dimensions via TSNE (Maaten and Hinton, 2008). Figure 2 shows 20 distinct clusters, one for each sentence. This indicates that distances between different vectors representing the same target sentence are much smaller than distances between vectors representing different sentences, and that distances in token space could be reflected in the latent steering space.

Motivated by the clustering results, we investigate whether the mean vector of the 8 extracted steering vectors for each target sentence recover the same sentence. Experiments show that mean vectors are able to recover target sentences nearly perfectly, leading to a BLEU-4 of 99.4, further establishing the robustness of our method.

#### 4 Is unsupervised style transfer in the latent steering space possible?

We explore whether vector arithmetic in this space is possible in the context of unsupervised style transfer. In other words, we measure whether adding an offset vector, which captures the desired style transfer, to the steering vector effectively changes the style of the generated sentence. Here, we show that unsupervised vector arithmetic with steering vectors is effective for unsupervised sentiment transfer, with performance comparable to models tailored to this task.

After extracting steering vectors for each sentence, we compute offset vectors by averaging steering vectors for a set of sentences in the source style  $\bar{z}_{source}$  and subtracting from the average of a set of steering vectors for the target style  $\bar{z}_{target}$ . Next, we flip the style of each sentence in our test set by adding the respective style transfer vector directly to its steering vector after scaling it by  $\lambda$ :

$$z_{totarget} = \bar{z}_{target} - \bar{z}_{source} \quad (2)$$

$$\hat{z}_{target} = z_{source} + \lambda \cdot z_{totarget} \quad (3)$$

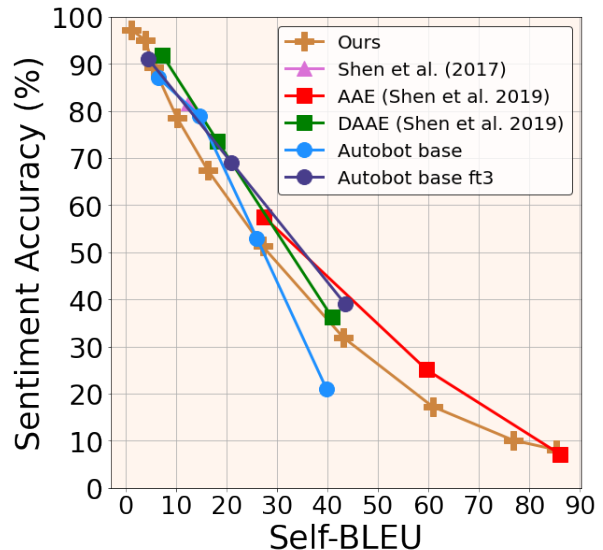


Figure 3: Evaluation of unsupervised sentiment transfer on the Yelp dataset. The plot shows accuracy vs. self-BLEU by varying  $\lambda = (0.25, 0.5, 1.0, 1.5, 2.0, 2.5, 3.0, 4.0, 5.0, 10.0)$  for our method. Overall, the steering vectors perform comparably to prior work.

**Unsupervised Sentiment Transfer** Using the Yelp Sentiment dataset preprocessed by Shen et al. (2017), we take 100 sentences from the validation set from each of the two classes of sentiment to compute offset vectors and evaluate on the test set. Following prior work (Shen et al., 2017), we measure how well this approach flips the sentiment of the sentence by measuring the accuracy of a RoBERTA-base model fine-tuned on the Yelp sentiment dataset. We also measure the BLEU-4 between the style transferred sentences and the original and report the results in Figure 3. We call this Self-BLEU following prior work. For this experiment, our steering vectors are injected after the 7th self-attention layer at the first timestep.

We find that simple vector arithmetic via our steering vectors, which is fully unsupervised, performs comparably to Shen et al. (2017), who learn an autoencoder-based model for the task in a fully supervised manner. Our method also compares well with the Autobot (Montero et al., 2021), AAE, and DAAE models (Shen et al., 2020), which although are unsupervised, either require training on in-domain data or require pretraining on millions of tokens in order to be effective. Other methods that use techniques from unsupervised machine translation to leverage the unpaired data in the task outperform all of these methods significantly (Hu et al., 2017; Lample et al., 2019; He et al., 2020).

Steering vectors	
Positive Input	the taste is excellent!
+0.5 * $z_{\text{negative}}$	the taste is excellent!
+1.0 * $z_{\text{negative}}$	the taste is excellent!
+1.5 * $z_{\text{negative}}$	the taste is bitter and bitter
+2.0 * $z_{\text{negative}}$	taste is bitter taste is bitter the taste is unpleasant.
Negative Input	the desserts were very bland.
+0.5 * $z_{\text{positive}}$	the desserts were very bland.
+1.0 * $z_{\text{positive}}$	the desserts were very bland.
+1.5 * $z_{\text{positive}}$	the desserts were very tasty.
+2.0 * $z_{\text{positive}}$	the desserts were very tasty.

Table 2: Examples of transferring sentiment using steering vectors for a positive input sentence (top) and negative input sentence (bottom). These results show fluency and accuracy in transfers while preserving the content of the input sentence.

These methods are not directly comparable to ours, as they evaluate on a different test set altogether and use the training set to train directly. Our method only requires access to 100 labeled examples per class to compute  $\bar{z}_{\text{source}}$  and  $\bar{z}_{\text{target}}$ , far fewer than other baselines. With as few as 10 examples per class, performance of our method remains competitive with autoencoder-based baselines.

Table 2 shows examples generated by our method for two input sentences. We find that resulting sentences become more positive or negative with increasing  $\lambda$  and often modify adjectives by swapping them out. On closer inspection, we find that fluency is often challenging for higher values of  $\lambda$  and that the generated sequences repeat individual words or phrases. In addition, we find that negative to positive sentiment transfer is qualitatively more fluent and accurate than positive to negative sentiment transfer; see Table 12 in the appendix for more example generations. Lastly, we evaluate on 19 paired style transfer tasks from the StylePTB dataset (Lyu et al., 2021), but modify the tasks to be unsupervised, following the same approach as above. We find that our method is similarly effective on these tasks; see Table 10 in the appendix for details.

## 5 Do distances between steering vectors reflect sentence similarity?

Previously, we found there exist multiple steering vectors that recover a target sentence and that those steering vectors are close together. This indicates the potential for distances in token space to be reflected in distances in the latent space occupied




Figure 4: On the test split of STS-B, we measure Spearman rank correlation ( $\rho \cdot 100$ ) between sentence similarity scores and cosine similarities between the steering vectors extracted from GPT2-117M when injected at different layers at the first timestep for those sentences. The vertical lines indicate extractive baselines: mean-pooled final hidden states for GPT2-117M and BERT-base as well as mean-pooled GloVe vectors. Results show that extracted steering vectors outperform these.

by steering vectors. In this section, we explore whether distances relate to semantic similarity. To do so, we use the STS-B test dataset, which consists of sentence pairs and similarity scores. To evaluate our method we extract steering vectors for each sentence separately, compute cosine similarity, and then correlate cosine similarity with annotator similarity via Spearman rank correlation.

In Figure 4, we show how well extracted steering vectors perform when injected at different layers and at the first timestep in the transformer stack. This observation mirrors the results from the experiment on recovery effectiveness: middle layers in the transformer stack are ideal for steering, leading to perfect recovery and highest performance on semantic similarity. We outperform mean pooling the final hidden states of GPT2-117M and BERT-base as well as averaged GloVe vectors. Even though our method is fully extractive, cosine distances reflect semantic similarity well. We take our two best performing configurations, the 7th self-attention layer and the 7th feedforward layer, and compare with unsupervised methods for tex-

Method	Spearman	Pearson
<i>Extractive methods</i>		
Avg GPT2-117M embeddings	25.92	16.52
Avg Bert embeddings	47.29	47.91
Avg GloVe embeddings	42.53	40.25
<b>Layer-7 self attention (ours)</b>	<b>52.04</b>	<b>51.17</b>
<b>Layer-7 feedforward (ours)</b>	<b>52.08</b>	<b>51.18</b>
<i>NLI-finetuned methods</i>		
AutoBot-base	58.49	-
InferSent - GloVe	68.03	-
<b>SBERT-NLI-base</b>	<b>77.03</b>	-
<i>Lexical methods</i>		
GloVe+UP	-	71.5
<b>GloVe+WR</b>	-	<b>72.0</b>

Table 3: We evaluate performance on the STS-B test set by measuring Spearman rank correlation and Pearson correlation ( $\rho \cdot 100$ ). We take our two best performing configurations from Figure 4 and compare them with three classes of unsupervised methods: extractive, NLI-finetuned, and lexical methods. Our method outperforms the extractive methods, but performs worse than the other methods, which are tailored for this task.

tual similarity. Table 3 shows that our extracted steering vectors out-perform prior extractive unsupervised methods. Predictably, however, methods which pretrain or fine-tune models on natural language inference datasets such as AutoBot (Montero et al., 2021), InferSent (Conneau et al., 2017), and SBERT (Reimers and Gurevych, 2019) perform better. Lexical methods tailored for semantic similarity such as GloVe with uSIF-weighting and piecewise component removal (GloVe + UP; Ethayarajh (2018)) and GloVe + WR (Arora et al., 2017) also outperform our method.

## 6 Analysis of Properties

### 6.1 Interpolation

Previous experiments indicate that the latent space occupied by steering vectors could be well-formed and smooth. To evaluate this qualitatively, we show linear interpolations of two pairs of steering vectors extracted from the Yelp Sentiment dataset in Figure 5. The space between the vectors look smooth with well-formed grammatical sentences that mix the content of two sentences effectively. The first interpolation (sentence pair 1) in Figure 5 shows that the positive sentiment of the first sentence car-

$\lambda$	Sentence Pair 1	Sentence Pair 2
0.0	all was hot and good.	four peaks also offers great food too.
0.1	all was hot and good.	four peaks also offers a great range of food and food products too.
0.2	all was hot and good.	four peaks also offers a number of other great food and drink options too.
0.3	all was hot and good.	four years ago
0.4	all was good	four years ago
0.5	all was good	when you're young.
0.6	when it came to my turkey sandwich, my husband and I had a great time.	when you are young, you simply cannot wait any longer.
0.7	when it came to my turkey sandwich, my turkey sandwich was pretty darn good.	when you are young, you simply do not know any better.
0.8	when it finally came, my turkey sandwich was kinda blah.	when you are that young, you simply do not know any better.
0.9	when it finally came, my turkey sandwich was kinda blah.	when you are that young, you simply don't know any better.
1.0	when it finally came, my turkey sandwich was kinda blah.	when you are that young, you simply don't know any better.

Figure 5: Interpolation between steering vectors extracted from two pairs of random sentences from the Yelp Sentiment test set. Decoding from interpolated vectors from two sentences produces well-formed output that incrementally changes the sentiment and meaning.

ries all the way to  $\lambda = 0.7$ , despite the content of the sentence changing to the second sentence. The second interpolation (sentence pair 2) in Figure 5 indicates that the latent space could encode some semantics relating to time. The second sentence includes the word "young" and so the transition between the two in  $\lambda = 0.3, 0.4$  combines the word "four" from the first sentence with the temporal component of "years ago" to relate the two sentences. Lastly, for each individual sentence there exists a radius around it where those vectors also steer the language model to generate the same target sentence. This could indicate that sentences have a representative volume from which, if any vector was sampled, could recover the sentence.

### 6.2 Sampling

Previous experiments show distances reflect semantic similarity and hint at the possibility that the latent space is smooth. Given this, we evaluate whether we can sample from this space. We take 4000 extracted steering vectors from the Yelp Sentiment test set. We treat each dimension of the steering vector as an independent random variable that is normally distributed with a mean and variance equal to the mean and variance across that dimension over this set of steering vectors. Table 11 shows the results of sampling 24 steering vectors and generating from them. We observe

mixed results: 5 samples lead to fully-formed sentences, and the remaining 19 lead to single tokens or phrases, indicating that treating steering vectors as samples from a multivariate Gaussian is not a reliable approach for sampling well-formed text.

### 6.3 Intrinsic Dimension & Space Complexity

We define the intrinsic dimension of the task of steering a language model as the minimum dimension of  $z_{steer}$  that achieves perfect recovery on a set of sentences. To measure intrinsic dimension, we vary the dimensions of  $z_{steer}$ , choosing 192, 384, 576, 768. We observe that reconstruction BLEU increases as the steering vector dimension increases, indicating that 768 dimensions may be needed to recover sequences nearly perfectly. Given this, we conclude that the intrinsic dimension is at most 768. However, a lower-dimensional representation can recover most sentences: 384 dimensions led to a reconstruction BLEU of 83.29. See Table 4 for more details. Additionally, we find that sentence length and reconstruction BLEU are inversely correlated, i.e. longer sequences are harder to recover. This is well-known; the number of bits needed to encode a sequence grows linearly with its length. We find that all four dimensions of steering vectors can recover short sentences, but lower dimensional steering vectors struggle to recover longer ones.

Steering vector dimension	192	384	576	768
Reconstruction BLEU-4	43.43	83.29	93.93	100.00

Table 4: Reconstruction BLEU for different steering vector dimensions. Sentence recovery increases monotonically as the dimension increases, up to 100% recovery at the model’s hidden dimension.

Since steering vectors do not depend on sequence length, space complexity may not be a problem. For a sequence of length 128, assuming 7 characters per word on average (including spaces), storage as a string takes  $128 * 7 = 896$  bytes. Our 768d steering vector uses 1536 bytes (fp16), but we can compress it by a factor of 2 (384d) sacrificing a little recovery (see Table 4) and store it using 768 bytes, less than its string representation.

### 6.4 Memorization

Our nearly perfect recoverability performance indicates that steering vectors could either be encoding important properties by leveraging the language




Figure 6: We measure reconstruction BLEU for steering vectors learned for three datasets: books, shuffled, and gibberish. Reconstruction BLEU for gibberish and shuffled data is lower than books indicating that the steering vector isn’t just memorizing the sequence, but also leveraging the language model well.

model, which would help generalization, or just simply be memorizing arbitrary sequences without using the underlying language model at all. In order to evaluate this, we randomly sample 64 sentences with lengths matching that of the books subset of our dataset, where each token is sampled uniformly at random with replacement from the vocabulary, and call this the gibberish fold of our dataset, following Subramani et al. (2019). Secondly, to measure whether both content and word order affect recoverability, we construct another fold, the shuffled fold, by randomly shuffling the tokens in the sentences in the books subset.

Figure 6 shows the results of injecting steering vectors into the 6th layer after the self-attention block in the transformer for all timesteps and the first timestep across all three datasets. We observe that recoverability is highest for books, then shuffled, and lastly gibberish. The gap between performance on books and gibberish indicates that steering vectors are not simply memorizing. Since recovery on books is greater than recovery on shuffled, we conclude that steering vectors encode some information about word order. Lastly, we notice that only passing the steering vector at the first timestep may reduce unwanted memorization capability because the relative difference in recovery between gibberish and the other sets is large.

### 6.5 Connection to Prompting

Motivated by the successes of prompt-based methods on zero-shot tasks with large generative lan-



guage models such as GPT-3 (Brown et al., 2020), we evaluate a prompt-based version of our method. Instead of adding  $z_{steer}$  to the hidden states of the language model, we concatenate  $k$  steering vectors with the input embeddings, so that all tokens can attend to these  $z_{steer}$  vectors. Experiments on the books subset show that recovery is much lower with this prompt-based approach than when injecting steering vectors directly into the transformer stack of the model. Even with  $k = 50$  steering vectors injected via this prompt-based approach, recovery fails to match that of a single steering vector  $z_{steer}$  injected into the hidden states of the language model.

Num prompt vectors	1	5	10	20	50
Reconstruction BLEU-4	81.7	94.3	98.7	98.6	98.5

Table 5: We measure reconstruction BLEU using a prompt-based approach, where latent steering vectors are concatenated to the embeddings. Even though each prompt vector is 768 dimensional, reconstruction BLEU is much lower in this setting than injecting a single steering vector into the layers of the transformer stack.

## 7 Related Work

There exist many works, often using text-based autoencoders that try to induce a sentence representation space for controllable text generation by learning new models (Hu et al., 2017; Shen et al., 2017, 2020; Mai et al., 2020; Montero et al., 2021). Our work concludes that we can extract steering vectors from pretrained models that have latent spaces that allow operations like this, without having to train any new models at all. Other approaches control language models by adapting their hidden states using steerable layers, adapters, or steering their logits using auxiliary language models (Gulcehre et al., 2015; Dathathri et al., 2019; Houlsby et al., 2019; Zhang et al., 2020; Liu et al., 2021; Krause et al., 2021). Our method differs from all of these: we extract steering vectors directly from a language model and operate on the latent space occupied by these vectors, never fine-tuning any component of the model. Subramani et al. (2019) investigate whether LSTM-based language models have sentence representations from which they can generate the original sentence. Although this premise relates to our first question: can we extract steering vectors, we extend far beyond that

showing that vector arithmetic in the context of unsupervised style transfer is effective in our latent steering space.

## 8 Conclusion

In this paper we introduce a different approach to controllable text generation, where we extract latent steering vectors directly from a pretrained language model without fine-tuning. Further, we find that our steering vectors lead to near perfect recovery on English sentences from a variety of domains. We show that vector arithmetic can be used in the context of unsupervised style transfer on the Yelp sentiment dataset and StylePTB benchmark, performing comparably to models tailored to these tasks. Experiments reveal that distances between steering vectors reflect sentence similarity when evaluated on STS-B, outperforming extractive methods. Finally, we analyze properties of the steering vectors. Our results indicate that we can control frozen pretrained language models effectively through their latent steering space.

## 9 Ethics Statement

We introduce a new approach for controllable text generation by extracting vectors from a pretrained language model, leveraging information that is already encoded in the language model. Large pretrained models are known to be biased and our method of extracting steering vectors can reflect biases already present in these large pretrained language models (Bender et al., 2021). The methods we present for controllable text generation could potentially be used for many downstream tasks such as unsupervised style transfer, abstractive summarization, and offensive content removal. Unfortunately, this also means that this technology has the potential to be misused to perpetuate biases or generate offensive or toxic text.

Our technology does not guarantee removal of toxic content, even in the case of unsupervised style transfer from toxic to nontoxic text. To use this method, we encourage readers to first take steps to address biases that are already present in the underlying language model. Further we recommend that this technology not be used in high-stakes settings, especially those where deployment of this technology could cause harm.

## References

- Sanjeev Arora, Yingyu Liang, and Tengyu Ma. 2017. A simple but tough-to-beat baseline for sentence embeddings. In *ICLR*.
- Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. On the dangers of stochastic parrots: Can language models be too big? . *FAccT*.
- Samuel R Bowman, Luke Vilnis, Oriol Vinyals, Andrew M Dai, Rafal Jozefowicz, and Samy Bengio. 2016. Generating sentences from a continuous space. *CoNLL 2016*.
- Tom B. Brown, Benjamin Pickman Mann, Nick Ryder, Melanie Subbiah, Jean Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, G. Krüger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric J 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. *arXiv preprint arXiv:2005.14165*.
- Isaac Caswell, Julia Kreutzer, Lisa Wang, Ahsan Wahaab, Daan van Esch, Nasanbayar Ulzii-Orshikh, Allahsera Auguste Tapo, Nishant Subramani, Artem Sokolov, Claytone Sikasote, Monang Setyawan, Supheakmongkol Sarin, Sokhar Samb, Benoît Sagot, Clara Rivera, Annette Rios Gonzales, Isabel Papadimitriou, Salomey Osei, Pedro Ortiz Suarez, Iroro Orife, Kelechi Ogueji, Rubungo Andre Niyongabo, Toan Q. Nguyen, Mathias Muller, Andr’e Muller, Shamsuddeen Hassan Muhammad, Nanda Firdausi Muhammad, Ayanda Mnyakeni, Jamshidbek Mirzakhlov, Tapiwanashe Matangira, Colin Leong, Nze Lawson, Sneha Kudugunta, Yacine Jernite, M. Jenny, Orhan Firat, Bonaventure F. P. Dossou, Sakhile Dlamini, Nisansa de Silva, Sakine cCabuk Balli, Stella Rose Biderman, Alessia Battisti, Ahmed Baruwa, Ankur Bapna, Pallavi N. Baljekar, Israel Abebe Azime, Ayodele Awokoya, Duygu Ataman, Orevaoghene Ahia, Oghenefego Ahia, Sweta Agrawal, and Mofetoluwa Adeyemi. 2022. Quality at a glance: An audit of web-crawled multilingual datasets. *TACL*.
- Daniel Cer, Mona Diab, Eneko Agirre, Iñigo Lopez-Gazpio, and Lucia Specia. 2017. [SemEval-2017 task 1: Semantic textual similarity multilingual and crosslingual focused evaluation](#). In *Proceedings of the 11th International Workshop on Semantic Evaluation (SemEval-2017)*, pages 1–14, Vancouver, Canada. Association for Computational Linguistics.
- Boxing Chen and Colin Cherry. 2014. A systematic comparison of smoothing techniques for sentence-level bleu. In *Proceedings of the Ninth Workshop on Statistical Machine Translation*.
- Alexis Conneau, Douwe Kiela, Holger Schwenk, Loïc Barrault, and Antoine Bordes. 2017. [Supervised learning of universal sentence representations from natural language inference data](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 670–680, Copenhagen, Denmark. Association for Computational Linguistics.
- Cristian Danescu-Niculescu-Mizil and Lillian Lee. 2011. Chameleons in imagined conversations: A new approach to understanding coordination of linguistic style in dialogs. In *CMCL@ACL*.
- Sumanth Dathathri, Andrea Madotto, Janice Lan, Jane Hung, Eric Frank, Piero Molino, Jason Yosinski, and Rosanne Liu. 2019. Plug and play language models: A simple approach to controlled text generation. In *ICLR*.
- Yuntian Deng, Anton Bakhtin, Myle Ott, and Arthur D. Szlam. 2020. Residual energy-based models for text generation. In *ICLR*.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2018. Bert: Pre-training of deep bidirectional transformers for language understanding. *CoRR*.
- Kawin Ethayarajh. 2018. [Unsupervised random walk sentence embeddings: A strong but simple baseline](#). In *Proceedings of The Third Workshop on Representation Learning for NLP*, pages 91–100, Melbourne, Australia. Association for Computational Linguistics.
- Kawin Ethayarajh. 2019. [How contextual are contextualized word representations? Comparing the geometry of BERT, ELMo, and GPT-2 embeddings](#). In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 55–65, Hong Kong, China. Association for Computational Linguistics.
- David Graff, Junbo Kong, Ke Chen, and Kazuaki Maeda. 2003. English gigaword. *Linguistic Data Consortium, Philadelphia*.
- Jiatao Gu, Kyunghyun Cho, and Victor O.K. Li. 2017. [Trainable greedy decoding for neural machine translation](#). In *Proceedings of the 2017 Conference on Empirical Methods in Natural Language Processing*, pages 1968–1978, Copenhagen, Denmark. Association for Computational Linguistics.
- Caglar Gulcehre, Orhan Firat, Kelvin Xu, Kyunghyun Cho, Loic Barrault, Hui-Chi Lin, Fethi Bougares, Holger Schwenk, and Yoshua Bengio. 2015. On using monolingual corpora in neural machine translation. *arXiv preprint arXiv:1503.03535*.
- Junxian He, Xinyi Wang, Graham Neubig, and Taylor Berg-Kirkpatrick. 2020. A probabilistic formulation of unsupervised text style transfer. In *ICLR*.

- Neil Houlsby, Andrei Giurgiu, Stanislaw Jastrzebski, Bruna Morrone, Quentin de Laroussilhe, Andrea Gesmundo, Mona Attariyan, and Sylvain Gelly. 2019. Parameter-efficient transfer learning for nlp. In *ICML*.
- Zhiting Hu, Zichao Yang, Xiaodan Liang, Ruslan Salakhutdinov, and Eric P. Xing. 2017. Toward controlled generation of text. In *ICML*.
- Di Jin, Zhijing Jin, Zhiting Hu, Olga Vechtomova, and Rada Mihalcea. 2020. Deep learning for text style transfer: A survey. *ArXiv*, abs/2011.00416.
- Yoon Kim. 2021. Sequence-to-sequence learning with latent neural grammars. *ArXiv*, abs/2109.01135.
- Diederik P Kingma and Jimmy Ba. 2014. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*.
- Ben Krause, Akhilesh Deepak Gotmare, Bryan McCann, Nitish Shirish Keskar, Shafiq Joty, Richard Socher, and Nazneen Fatema Rajani. 2021. **GeDi: Generative discriminator guided sequence generation**. In *Findings of the Association for Computational Linguistics: EMNLP 2021*, pages 4929–4952, Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Guillaume Lample, Sandeep Subramanian, Eric Michael Smith, Ludovic Denoyer, Marc’Aurelio Ranzato, and Y-Lan Boureau. 2019. Multiple-attribute text rewriting. In *ICLR*.
- Marie Lebert. 2008. Project gutenber (1971-2008).
- Brian Lester, Rami Al-Rfou, and Noah Constant. 2021. **The power of scale for parameter-efficient prompt tuning**. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 3045–3059, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Yongjie Lin, Yi Chern Tan, and Robert Frank. 2019. **Open sesame: Getting inside BERT’s linguistic knowledge**. In *Proceedings of the 2019 ACL Workshop BlackboxNLP: Analyzing and Interpreting Neural Networks for NLP*, pages 241–253, Florence, Italy. Association for Computational Linguistics.
- Alisa Liu, Maarten Sap, Ximing Lu, Swabha Swayamdipta, Chandra Bhagavatula, Noah A. Smith, and Yejin Choi. 2021. **DExperts: Decoding-time controlled text generation with experts and anti-experts**. In *Proceedings of the 59th Annual Meeting of the Association for Computational Linguistics and the 11th International Joint Conference on Natural Language Processing (Volume 1: Long Papers)*, pages 6691–6706, Online. Association for Computational Linguistics.
- Yiwei Lyu, Paul Pu Liang, Hai Pham, Eduard Hovy, Barnabás Póczos, Ruslan Salakhutdinov, and Louis-Philippe Morency. 2021. **StylePTB: A compositional benchmark for fine-grained controllable text style transfer**. In *Proceedings of the 2021 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 2116–2138, Online. Association for Computational Linguistics.
- L. V. D. Maaten and Geoffrey E. Hinton. 2008. Visualizing data using t-sne. *JMLR*.
- Florian Mai, Nikolaos Pappas, Ivan Montero, Noah A. Smith, and James Henderson. 2020. **Plug and play autoencoders for conditional text generation**. In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 6076–6092, Online. Association for Computational Linguistics.
- Stephen Merity, Caiming Xiong, James Bradbury, and Richard Socher. 2017. Pointer sentinel mixture models. *ArXiv*, abs/1609.07843.
- Ivan Montero, Nikolaos Pappas, and Noah A. Smith. 2021. **Sentence bottleneck autoencoders from transformer language models**. In *Proceedings of the 2021 Conference on Empirical Methods in Natural Language Processing*, pages 1822–1831, Online and Punta Cana, Dominican Republic. Association for Computational Linguistics.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. Bleu: a method for automatic evaluation of machine translation. In *ACL*.
- 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, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke S. Zettlemoyer. 2018. Deep contextualized word representations. In *NAACL-HLT*.
- Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. Improving language understanding by generative pre-training. Unpublished ms. available through a link at <https://blog.openai.com/language-unsupervised/>.
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. Language models are unsupervised multitask learners.
- Nils Reimers and Iryna Gurevych. 2019. **Sentence-BERT: Sentence embeddings using Siamese BERT-networks**. In *Proceedings of the 2019 Conference on Empirical Methods in Natural Language Processing and the 9th International Joint Conference on Natural Language Processing (EMNLP-IJCNLP)*, pages 3982–3992, Hong Kong, China. Association for Computational Linguistics.

- Anna Rogers, Olga Kovaleva, and Anna Rumshisky. 2020. [A primer in BERTology: What we know about how BERT works](#). *Transactions of the Association for Computational Linguistics*, 8:842–866.
- Alexis Ross, Tongshuang (Sherry) Wu, Hao Peng, Matthew E. Peters, and Matt Gardner. 2021. Tailor: Generating and perturbing text with semantic controls. *ArXiv*, abs/2107.07150.
- Tianxiao Shen, Tao Lei, Regina Barzilay, and T. Jaakkola. 2017. Style transfer from non-parallel text by cross-alignment. In *NIPS*.
- Tianxiao Shen, Jonas Mueller, Regina Barzilay, and T. Jaakkola. 2020. Educating text autoencoders: Latent representation guidance via denoising. In *ICML*.
- Taylor Shin, Yasaman Razeghi, Robert L. Logan IV, Eric Wallace, and Sameer Singh. 2020. [AutoPrompt: Eliciting Knowledge from Language Models with Automatically Generated Prompts](#). In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 4222–4235, Online. Association for Computational Linguistics.
- Nishant Subramani, Samuel Bowman, and Kyunghyun Cho. 2019. Can unconditional language models recover arbitrary sentences? In *NeurIPS*.
- 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 *NeurIPS*.
- Pascal Vincent, H. Larochelle, Yoshua Bengio, and Pierre-Antoine Manzagol. 2008. Extracting and composing robust features with denoising autoencoders. In *ICML*.
- Lilian Weng. 2021. [Controllable neural text generation](#). *lilianweng.github.io/lil-log*.
- Jeffrey O. Zhang, Alexander Sax, Amir Roshan Zamir, Leonidas J. Guibas, and Jitendra Malik. 2020. Side-tuning: A baseline for network adaptation via additive side networks. In *ECCV*.
- Xiang Zhang, Junbo Zhao, and Yann LeCun. 2015. Character-level convolutional networks for text classification. *NIPS*.

## A Appendix

### A.1 Extracting steering vectors

In this section, we show the hyperparameter configurations used for extracting steering vectors from GPT2-117M. Table 6 contains the list of final hyperparameters that we use to extract steering vectors for the different analyses in this paper. Table 7 shows the recovery performance of steering vectors when injected at different layers in the transformer stack on our compiled dataset. These experiments reveal that injecting in the middle of the transformer stack either after the self attention layer or the feedforward layer leads to the highest BLEU-4 performance. In fact, any layer other than the first or last layer achieves nearly perfect recovery.

In Table 8 we look at recovery performance when injecting steering vectors at the embedding layer, transformer stack, and language modeling head, as well as different combinations of them. Injecting steering vectors at every layer in the transformer stack performed best. Table 9 shows how recoverability changes with respect to how many timesteps  $z_{steer}$  is injected at. Injecting at all timesteps performs negligibly better than injecting at just the first timestep.

Hyperparameters	Values
Model	GPT-2-117M
Max train steps	500
Vector initialization strategy	Xavier normal
Learning rate	[0.01, 1.0]
Optimizer	Adam
Learning rate Scheduler	Decay on a plateau
Scheduler decay factor	0.9
Scheduler decay patience	1.0

Table 6: List of hyperparameter configurations used to extract  $z_{steer}$  from GPT2-117M.

### A.2 Unsupervised Sentiment Transfer

**Yelp Sentiment** We also include generations from the unsupervised sentiment transfer experiment on the Yelp dataset. Table 12 shows 8 more generations. These generations highlight the same trends as before: with increasing  $\lambda$ , sentiment transfer strength increases. We find that some genera-

Injection location	layers	timestep	lr	BLEU-4
self_attn	0	all timesteps	1	33.25
feedforward	0	all timesteps	1	97.68
self_attn	1	all timesteps	1	98.06
feedforward	1	all timesteps	1	99.54
self_attn	2	all timesteps	1	100.00
feedforward	2	all timesteps	1	99.69
self_attn	3	all timesteps	1	100.00
feedforward	3	all timesteps	1	100.00
self_attn	4	all timesteps	1	100.00
feedforward	4	all timesteps	1	100.00
self_attn	5	all timesteps	1	100.00
feedforward	5	all timesteps	1	100.00
self_attn	6	all timesteps	1	<b>100.00</b>
feedforward	6	all timesteps	1	99.62
self_attn	7	all timesteps	1	99.62
feedforward	7	all timesteps	1	<b>100.00</b>
self_attn	8	all timesteps	1	100.00
feedforward	8	all timesteps	1	98.84
self_attn	9	all timesteps	1	99.22
feedforward	9	all timesteps	1	98.61
self_attn	10	all timesteps	1	97.50
feedforward	10	all timesteps	1	95.24
self_attn	11	all timesteps	1	86.04
feedforward	11	all timesteps	1	6.29

Table 7: This table shows the reconstruction BLEU-4 for steering vectors from our compiled dataset when injected after different self attention and feedforward layers in the transformer stack. Injecting at the middle layer of the language model performs best.

tions do more than just flip the sentiment of the major adjective in the sentence such as adding the phrase "a great way to get a good laugh" in the 4th negative to positive generation when  $\lambda = 2.5$ .

**StylePTB** For this study, we use 19 of 21 paired style transfer tasks from the StylePTB dataset (Lyu et al., 2021), but modify the tasks to be unsupervised, following the same approach as sentiment transfer. We randomly sample 100 sentences for each class from the training split for each of the style classes and use those to compute offset vectors. This offset vector is then added to the steering vector of the sentence to transfer style. We follow the evaluation in Lyu et al. (2021) because we have ground truth data and compare with fully

<b>Injection location</b>	<b>timestep</b>	<b>lr</b>	<b>BLEU-4</b>
embedding	all timesteps	0.01	33.99
every_layer	all timesteps	0.01	100.00
lm_head	all timesteps	0.01	6.72
embedding+every_layer	all timesteps	0.01	96.52
every_layer+lm_head	all timesteps	0.01	100.00
embedding+lm_head	all timesteps	0.01	83.27
embedding+every_layer+lm_head	all timesteps	0.01	98.11
every_layer_self_attn	all timesteps	0.01	99.62
<b>every_layer+every_layer_self_attn</b>	<b>all timesteps</b>	<b>0.01</b>	<b>100.00</b>
every_layer_self_attn+embedding+lm_head	all timesteps	0.01	97.31
every_layer_self_attn+lm_head	all timesteps	0.01	99.62
every_layer_self_attn+embedding	all timesteps	0.01	94.28

Table 8: Here, we present the reconstruction BLEU-4 results for steering vectors on our multi-domain compiled dataset. We vary injection location here and observe that injecting into the transformer stack is necessary for good recovery. Injecting at the embedding or language model head performs poorly.

<b>Injection location</b>	<b>timestep</b>	<b>lr</b>	<b>BLEU-4</b>
every_layer+every_layer_self_attn	all timesteps	0.01	100.0
every_layer+every_layer_self_attn	first timestep	0.01	91.7
Layer 7 (feedforward)	all timesteps	1	100.0
Layer 7 (feedforward)	first timestep	1	99.2
Layer 6 (self_attn)	all timesteps	1	100.0
Layer 6 (self_attn)	first timestep	1	99.8

Table 9: In this table, we vary the timestep where we inject  $z_{steer}$  (all timesteps or first timestep) for three of our best injection locations. We again evaluate on our multi-domain compiled dataset and find that injecting at just the first timestep has a negligible decrease in recovery performance.

supervised methods. Experiments show that unsupervised vector arithmetic with steering vectors performs comparably using BLEU-1 to supervised methods designed for style transfer on tasks that require minimal edits (adjective emphasis (AEM), active to passive (ATP), information addition (IAD), and PP front to back (PFB)). We report BLEU-1 following prior work. See Table 10 for results on all 19 tasks. Note Lyu et al. (2021) do not report any baseline numbers for AAR, ASR, LFS, MFS, NAR, NSR, and VSR for any of their models.

### A.3 Sampling

In order to evaluate whether we can sample steering vectors reliably, we collect 4,000 extracted steering vectors from the Yelp Sentiment test set. To generate, we consider each dimension of the steering vector as an independent random variable that is normally distributed. The dimension means and

variances are equal to the mean and variance for that dimension across this set of steering vectors. In Table 11, we show the results of sampling 24 steering vectors from these independent normally distributed random variables and generating from them using GPT2-117M as our language model. These results are mixed with approximately 20% of the generations leading to fully formed sentences and the remaining 80% corresponding to individual words or short phrases. This could perhaps be partially explained by the fact that text from the web, including the corpora GPT2 was trained on, can often be of poor quality, especially when automatically crawled (Caswell et al., 2022). Alternatively, our choice of considering  $d$ -dimensional steering vectors as samples from  $d$  independent normally distributed random variables could be an incorrect assumption. Alternative formulations could lead to more fluent and reliable generations.

	Ours: $\lambda = 0.25$	GPT2-finetune	Seq2seq	TAILOR	Neural QCFG + copy	Retrieve-Edit
AAR	<b>0.825</b>	-	-	-	-	-
AEM	<b>0.774</b>	0.263	0.187	-	0.676	0.387
ARR	0.721	0.647	0.450	0.781	-	<b>0.897</b>
ASR	<b>0.819</b>	-	-	-	-	-
ATP	0.666	0.476	0.373	0.556	<b>0.836</b>	0.681
IAD	<b>0.772</b>	0.479	0.345	-	-	0.493
LFS	<b>0.396</b>	-	-	-	-	-
MFS	<b>0.748</b>	-	-	-	-	-
NAR	<b>0.825</b>	-	-	-	-	-
NSR	<b>0.677</b>	-	-	-	-	-
PFB	0.819	0.398	0.393	<b>0.842</b>	-	0.541
PPR	0.393	0.763	0.330	0.717	-	<b>0.798</b>
PTA	0.574	0.433	0.339	-	-	<b>0.714</b>
SBR	0.120	0.430	0.317	-	-	<b>0.706</b>
TFU	0.699	0.895	0.527	0.873	-	<b>0.899</b>
TPA	0.478	0.836	0.478	0.884	-	<b>0.935</b>
TPR	0.692	0.754	0.516	0.710	-	<b>0.909</b>
VEM	0.548	0.309	0.289	-	<b>0.664</b>	0.416
VSR	<b>0.739</b>	-	-	-	-	-

Table 10: In this table, we show performance on StylePTB. Although our method is unsupervised, we outperform GPT2-finetune and seq2seq on most tasks. For minimal edit tasks such as AEM, ARR, ATP, and PFB, we achieve comparable performance to TAILOR, Neural QCFG + copy, and Retrieve-Edit, which are models trained specifically for these types of tasks. Note: we obtain the numbers for GPT2-finetune, Seq2seq, and Retrieve-Edit from (Lyu et al., 2021), for TAILOR from (Ross et al., 2021), and for Neural QCFG+copy from (Kim, 2021).

Sampled Sequences	
...	mobile
wine..	the first time that we've seen a team that looked good on paper.
peopled by.	Gathering around the world, we can all agree that the next step is to get our voices heard.
kitchen.....	x
life	item link
nomnomnomnom	appointments
of	kitate.com
We're going to make sure that we have a safe and secure environment for our employees.	3
app	hotel
racial	imagine a world where every day we see a new voice in our communities.
applify	(AAP) - The United States and its European allies are pressing ahead with plans to boost the number of refugees arriving in the country from Iraq and Syria.\n\nThe United States and its allies are pressing ahead on the issue as they work to boost the number and scope of refugees arriving in Europe.
iv	the best.

Table 11: Here we show results from our sampling experiment, where we treat steering vectors as samples from  $d$  independent normally distributed random variables. We sample 24 steering vectors, pass them to GPT2-117M, and decode, resulting in the 24 generations presented here.

Unsupervised sentiment transfer using steering vectors			
Positive to negative		Negative to positive	
input	i highly recommend this place!	input	my goodness it was so gross.
+0.5 * $Z_{tonegative}$	i highly recommend this place!	+0.5 * $Z_{topositive}$	my goodness it was so gross.
+1.0 * $Z_{tonegative}$	i highly recommend this place!	+1.0 * $Z_{topositive}$	my goodness it was so gross.
+1.5 * $Z_{tonegative}$	i highly recommend this place!	+1.5 * $Z_{topositive}$	my goodness it was so gross.
+2.0 * $Z_{tonegative}$	i was very disappointed.	+2.0 * $Z_{topositive}$	my goodness it was so good.
input	it is always good to find quality local spots when traveling.	input	went here for the first time to try something new ... bad idea.
+0.5 * $Z_{tonegative}$	it is always good to find quality local spots when traveling.	+0.5 * $Z_{topositive}$	went here for the first time to try something new.
+1.0 * $Z_{tonegative}$	it is always good to find quality local spots when traveling.	+1.0 * $Z_{topositive}$	went here for the first time to try something new.
+1.5 * $Z_{tonegative}$	it is always good to find local spots when traveling.	+1.5 * $Z_{topositive}$	went here for the first time to try something new.
+2.0 * $Z_{tonegative}$	it was always going to be a long time.	+2.0 * $Z_{topositive}$	went here for the first time to try something new. I'm really looking forward to trying something new for the first time.
input	it was delicious!	input	if i could give them a zero star review i would!
+0.5 * $Z_{tonegative}$	it was delicious!	+0.5 * $Z_{topositive}$	if i could give them a star i would!
+1.0 * $Z_{tonegative}$	it was delicious!	+1.0 * $Z_{topositive}$	if i could give them a star i would!
+1.5 * $Z_{tonegative}$	it was a very bad night.	+1.5 * $Z_{topositive}$	if i could give them a star i would!
+2.0 * $Z_{tonegative}$	it was a very bad night.	+2.0 * $Z_{topositive}$	if i could give them a star i would!
input	the food is fresh and the environment is good.	input	fries are n't worth coming back.
+0.5 * $Z_{tonegative}$	the food is fresh and the environment is good.	+0.5 * $Z_{topositive}$	fries are good.
+1.0 * $Z_{tonegative}$	the food is fresh and the environment is good.	+1.0 * $Z_{topositive}$	fries are good.
+1.5 * $Z_{tonegative}$	the food is fresh and the environment is good.	+1.5 * $Z_{topositive}$	fries are good.
+2.0 * $Z_{tonegative}$	the food is bad.	+2.0 * $Z_{topositive}$	fries are good.
+2.5 * $Z_{tonegative}$	the food was produced in the past.	+2.5 * $Z_{topositive}$	fries are a great way to get a good laugh.

Table 12: This table shows some generations from unsupervised sentiment transfer of steering vectors. Sentences are from the Yelp dataset. We find that with increasing  $\lambda$  sentiment transfers more strongly towards positive or negative, often switching at  $\lambda = 1.5$ .