

Generating Sentences by Editing Prototypes

Kelvin Guu*² Tatsunori B. Hashimoto*^{1,2} Yonatan Oren² Percy Liang^{1,2}
(* equal contribution)

¹Department of Computer Science ²Department of Statistics
Stanford University

{kguu, thashim, yonatano}@stanford.edu pliang@cs.stanford.edu

Abstract

We propose a new generative model of sentences that first samples a prototype sentence from the training corpus and then edits it into a new sentence. Compared to traditional models that generate from scratch either left-to-right or by first sampling a latent sentence vector, our prototype-then-edit model improves perplexity on language modeling and generates higher quality outputs according to human evaluation. Furthermore, the model gives rise to a latent edit vector that better captures interpretable semantics such as sentence similarity and analogies.

1 Introduction

The ability to generate sentences is a core component of many NLP systems, including speech recognition (Jurafsky and Martin, 2000), machine translation (Kalchbrenner and Blunsom, 2013; Koehn et al., 2007), summarization (Hahn and Mani, 2000; Nallapati et al., 2016), and dialogue (Ritter et al., 2011). State-of-the-art models are largely based on recurrent neural language models (NLMs) that generate sentences from scratch, often in a left-to-right manner. It is often observed that such NLMs suffer from the problem of favoring generic utterances such as “I don’t know” (Li et al., 2016). At the same time, common strategies to increase diversity often compromise grammaticality (Shao et al., 2017), suggesting that current NLMs simply lack the inductive bias to faithfully represent the full diversity of complex utterances.

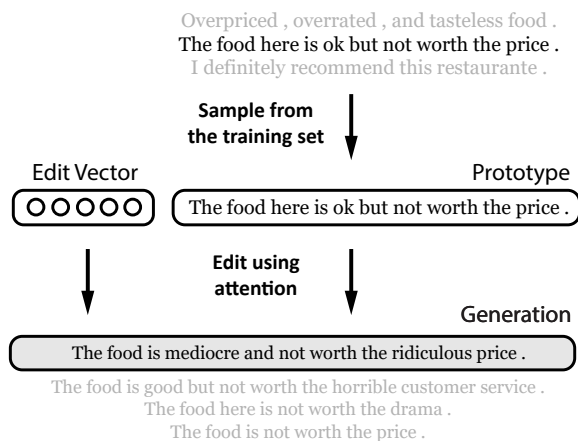


Figure 1: The prototype-then-edit model generates a sentence by sampling a random example from the training set and then editing it using a randomly sampled edit vector.

Indeed, it is difficult even for humans to write complex text from scratch in a single pass; we often create an initial draft and incrementally revise it (Hayes and Flower, 1986). Inspired by this process, we propose a new generative model of text which we call the prototype-then-edit model, illustrated in Figure 1. It first samples a random *prototype* sentence from the training corpus, and then invokes a *neural editor*, which draws a random “edit vector” and generates a new sentence by attending to the prototype while conditioning on the edit vector. The motivation is that sentences from the corpus provide a high quality starting point: they are grammatical, naturally diverse, and exhibit no bias towards shortness or vagueness. The attention mechanism (Bahdanau et al., 2015) then extracts the rich information from the prototype and generalizes to novel sentences.

We train the neural editor by maximizing

an approximation to the generative model’s log-likelihood. This objective is a sum over lexically-similar sentence pairs in the training set, which we can scalably approximate using locality sensitive hashing. We also show empirically that most lexically similar sentences are also semantically similar, thereby endowing the neural editor with additional semantic structure. For example, we can use the neural editor to perform a random walk from a seed sentence to traverse semantic space.

We compare our prototype-then-edit model against approaches which generate from scratch on two fronts: language generation quality and semantic properties. For the former, our model generates higher quality generations according to human evaluations, and improves perplexity by 13 points on the Yelp corpus and 7 points on the One Billion Word Benchmark. For the latter, we show that the edit vectors outperform standard sentence variational autoencoders (Bowman et al., 2016) on semantic similarity, locally-controlled text generation, and sentence analogies tasks.

2 Problem statement

Our primary goal is to learn a generative model of sentences. In particular, we model sentence generation as a prototype-then-edit process:¹

1. **Prototype selector:** Given a training corpus of sentences \mathcal{X} , we randomly sample a *prototype* sentence, $x' \sim p(x')$ (in our case, uniform over \mathcal{X}).
2. **Neural editor:** First we draw an edit vector z , from an edit distribution $p(z)$, which encodes the type of edit. Then, we draw the final sentence x from $p_{\text{edit}}(x | x', z)$, which takes the prototype x' and the edit vector z .

Under this model, the likelihood of a sentence is:

$$p(x) = \sum_{x' \in \mathcal{X}} p(x | x')p(x') \quad (1)$$

$$p(x | x') = \int_z p_{\text{edit}}(x | x', z)p(z)dz, \quad (2)$$

¹For many applications such as machine translation or dialogue generation, there is a context (e.g. foreign sentence, dialogue history), which can be supplied to both the prototype selector and the neural editor. This paper focuses on the unconditional case.

where both prototype x' and edit vector z are latent variables.

Our formulation stems from the observation that many sentences in a large corpus can be represented as minor transformations of other sentences. For example, in the Yelp dataset (Yelp, 2017) we find that 70% of the test set is within Jaccard distance 0.5 of a training set sentence, even though almost no sentences are repeated verbatim. This implies that a neural editor which models lexically similar sentences should be an effective generative model for large parts of the test set.

A secondary goal for the neural editor is to capture certain semantic properties; we focus on the following two in particular:

1. **Semantic smoothness:** an edit should be able to alter the semantics of a sentence by a small and well-controlled amount, while multiple edits should make it possible to accumulate a larger change.
2. **Consistent edit behavior:** the edit vector z should model/control the variation in the type of edit that is performed. When we apply the same edit vector on different sentences, the neural editor should perform *semantically analogous* edits across the sentences.

In Section 4, we show that the neural editor can successfully capture both properties based on human evaluation.

3 Approach

We would like to train our neural editor by maximizing the marginal likelihood (Equation 1), but the exact likelihood is difficult to maximize because it involves a sum over all latent prototypes x' (expensive) and integration over the latent edit vector z (intractable).

We therefore propose two approximations to overcome these challenges:

1. We approximate the sum over all prototypes x' by only summing over x' that are *lexically similar* to x .
2. We lower bound the integral over latent edit vectors by modeling z with a variational autoencoder, which admits tractable inference via

the evidence lower bound (ELBO), which incidentally also provides additional semantic structure.

We describe and motivate both of these approximations in the subsections below.

3.1 Approximate sum on prototypes, x'

Equation 1 defines the probability of generating a sentence x as the total probability of reaching x via edits from every prototype $x' \in \mathcal{X}$. However, most prototypes are unrelated and should have very small probability of transforming into x . Therefore, we approximate the marginal distribution over prototypes by only considering the prototypes x' with high lexical overlap with x , as measured by word token Jaccard distance d_J . Formally, define a lexical similarity neighborhood as $\mathcal{N}(x) = \{x' \in \mathcal{X} : d_J(x, x') < 0.5\}$. The neighborhoods $\mathcal{N}(x)$ can be constructed efficiently with locality sensitive hashing (LSH) and minhashing. The full procedure is described in Appendix 6.1.

Then the log-likelihood of the prototype-then-edit process is lower bounded by:

$$\begin{aligned} \log p(x) &\geq \log \left(\sum_{x' \in \mathcal{N}(x)} p(x') p(x | x') \right) \\ &\geq \log \left(\sum_{x' \in \mathcal{N}(x)} |\mathcal{N}(x)|^{-1} p(x | x') \right) - \log |\mathcal{X}| \\ &\geq |\mathcal{N}(x)|^{-1} \sum_{x' \in \mathcal{N}(x)} \log p(x | x') - \log |\mathcal{X}| \quad (3) \end{aligned}$$

where the inequalities follow from summing over fewer terms, multiplying by a number that is at most 1, and Jensen’s inequality. Treating $|\mathcal{N}(x)|^{-1}$ as a constant² and summing over the training set, we arrive at the following objective:

$$\mathcal{L}_{\text{Lex}} = \sum_{x \in \mathcal{X}} \sum_{x' \in \mathcal{N}(x)} \log p(x | x'). \quad (4)$$

Interlude: lexical similarity semantics. We have motivated lexical similarity neighborhoods via computational considerations, but we find that lexical similarity training also captures semantic similarity. One can certainly construct sentences with small lexical distance that differ semantically (e.g., insertion of the word “not”). However, since we mine

²In preliminary experiments, we found this to yield better results.

sentences from a corpus grounded in real world events, most lexically similar sentences are also semantically similar. For example, given “my son enjoyed the delicious pizza”, we are far more likely to see “my son enjoyed the delicious macaroni”, versus “my son hated the delicious pizza”.

Human evaluations of 250 edit pairs sampled from lexical similarity neighborhoods on the Yelp corpus support this conclusion. 35.2% of the sentence pairs were judged to be exact paraphrases, while 84% of the pairs were judged to be at least roughly equivalent. Sentence pairs were negated or change in topic only 7.2% of the time. Thus, a neural editor trained on this distribution should preferentially generate semantically similar edits.

3.2 Approximate integration on edit vectors, z

Now let us tackle integration over the latent edit vectors. To do this, let us introduce the evidence lower bound (ELBO) to the integral over z in Equation 1:

$$\begin{aligned} \log p(x | x') &= \log \int_z p_{\text{edit}}(x | x', z) p(z) dz \\ &\geq \ell(x, x') = \mathbb{E}_{z \sim q(z | x, x')} [\log p_{\text{edit}}(x | x', z)] \\ &\quad - D_{\text{KL}}(q(z | x, x') || p(z)). \quad (5) \end{aligned}$$

The important elements of $\ell(x, x')$ are the *neural editor* $p_{\text{edit}}(x | x', z)$, the *edit prior* $p(z)$, and the *approximate edit posterior* $q(z | x, x')$ (we describe each of these shortly). Combining the ELBO with Equation 3, our final objective function is:

$$\mathcal{L}_{\text{Lex}} \geq \mathcal{L}_{\text{ELBO}} = \sum_{x \in \mathcal{X}} \sum_{x' \in \mathcal{N}(x)} \ell(x, x') \quad (6)$$

$\mathcal{L}_{\text{ELBO}}$ is now maximized over the parameters of both p_{edit} and q . Note that q is only used for training, and discarded at test time. With the introduction of q , we are now training $p(x | x')$ as a conditional variational autoencoder (C-VAE) (it is *conditioned* on x'). We maximize the objective via SGD, approximating the gradient using the usual Monte Carlo estimate for VAEs (Kingma and Welling, 2014).

Neural editor $p_{\text{edit}}(x | x', z)$: We implement our neural editor as a left-to-right sequence-to-sequence model with attention, where the prototype is the input sequence and the revised sentence is the output sequence. We employ an encoder-decoder architecture similar to Wu (2016), extending it to condition

on an edit vector z by concatenating z to the input of the decoder at each time step. Further details are given in Appendix 6.2.

Edit prior $p(z)$: We sample z from the prior by drawing a random magnitude $z_{\text{norm}} \sim \text{Unif}(0, 10)$ and then drawing a random direction $z_{\text{dir}} \sim \text{vMF}(0)$, where $\text{vMF}(0)$ is a uniform distribution over the unit sphere (von-Mises Fisher distribution with concentration = 0). The resulting $z = z_{\text{norm}}z_{\text{dir}}$.

Approximate edit posterior $q(z \mid x, x')$: q is named the *approximate edit posterior* because the ELBO is tight when q matches the true posterior: the best possible estimate of the edit z given both the prototype x' and the revision x . Our design of q treats the edit vector z as a generalization of word vectors. In the case of a single word insertion, a good edit vector would be the word vector of the inserted word. Extending this to multi-word edits, we would like multi-word insertions to be represented as the sum of the individual word vectors.

Since q is an encoder observing both the prototype x' and the revised sentence x , it can directly observe the word differences between x and x' . Define $I = x \setminus x'$ to be the set of words added to the prototype, and $D = x' \setminus x$ to be the words deleted.

We would then like q to output a z equal to

$$f(x, x') = \sum_{w \in I} \Phi(w) \oplus \sum_{w \in D} \Phi(w).$$

where $\Phi(w)$ is the word vector for word w and \oplus denotes concatenation. The word embeddings Φ are parameters of q .

However, q cannot deterministically output $f(x, x')$ — without any entropy in q , the KL term in equation 5 would be infinity and training would be infeasible. Hence, we design q to output a noisy, perturbed version of f : we perturb the norm of f by adding uniform noise, and we perturb the direction of f by adding von-Mises Fisher noise. The von-Mises Fisher distribution is a distribution over vectors with unit norm, with a mean μ and a precision κ such that the log-likelihood of drawing a vector decays linearly with the cosine similarity to μ .

Let $f_{\text{norm}} = \|f\|$ and $f_{\text{dir}} = f/f_{\text{norm}}$. Furthermore, define $\tilde{f}_{\text{norm}} = \min(f_{\text{norm}}, 10)$ to be the trun-

cated norm. Then,

$$q(z_{\text{dir}} \mid x, x') = \text{vMF}(z_{\text{dir}}; f_{\text{dir}}, \kappa) \propto \exp(\kappa z_{\text{dir}}^\top f_{\text{dir}})$$

$$q(z_{\text{norm}} \mid x, x') = \text{Unif}(z_{\text{norm}}; [\tilde{f}_{\text{norm}}, \tilde{f}_{\text{norm}} + \epsilon]),$$

where the resulting $z = z_{\text{dir}}z_{\text{norm}}$. The resulting distribution q has three parameters: (Φ, κ, ϵ) , where κ, ϵ are hyperparameters. This distribution is straightforward to use as part of a variational autoencoder, as sampling can be easily performed using the reparameterization trick and the rejection sampler of Wood (1994). Furthermore, the KL term has a closed form expression independent of z .

$$D_{\text{KL}}(\text{vMF}(\kappa) \parallel \text{vMF}(0)) = \kappa \frac{I_{d/2+1}(\kappa) + I_{d/2}(\kappa) \frac{d}{2\kappa}}{I_{d/2}(\kappa) - \frac{d}{2\kappa}} + \frac{d}{2} \log(\kappa/2) - \log(I_{d/2}(\kappa)\Gamma(d/2 + 1)),$$

where $I_n(\kappa)$ is the modified Bessel function of the first kind, and Γ is the gamma function.

Our design of q differs substantially from the standard choice of a standard normal distribution with a given mean (Bowman et al., 2016; Kingma and Welling, 2014) for two reasons:

First, by construction, edit vectors are sums of word vectors and since cosine distances are traditionally used to measure distances between word vectors, it would be natural to encode distances between edit vectors by the cosine distance. The von-Mises Fisher distribution captures this idea, as the log likelihood of transforming f_{dir} into z_{dir} decays linearly with the cosine similarity.

Second, our parameterization avoids collapsing the latent code, which is a serious problem with variational autoencoders in practice (Bowman et al., 2016). With a Gaussian latent noise variable, the KL-divergence term is instance-dependent, and thus the encoder must decide how much information to pass to the decoder. Even with training techniques such as annealing, the model often learns to ignore the encoder entirely. In our case, this does not occur, since the KL divergence between a von-Mises Fisher with parameter κ and the uniform distribution is independent of the mean f_{dir} , allowing us to optimize κ separately using binary search. In practice, we never observe issues with encoder collapse using standard gradient training.

4 Experiments

We divide our experimental results into two parts. In Section 4.3, we evaluate the merits of the prototype-then-edit model as a generative modeling strategy, measuring its improvements on language modeling (perplexity) and generation quality (human evaluations of diversity and plausibility). In Section 4.4, we focus on the semantics learned by the model and its latent edit vector space. We demonstrate that it possesses interpretable semantics, enabling us to smoothly control the magnitude of edits, incrementally optimize sentences for target properties, and perform analogy-style sentence transformations.

4.1 Datasets

We evaluate perplexity on the Yelp review corpus (YELP, Yelp (2017)) and the One Billion Word Language Model Benchmark (BILLIONWORD, Chelba (2013)). For qualitative evaluations of generation quality and semantics, we focus on YELP as our primary test case, as we found that human judgments of semantic similarity were much better calibrated in this focused setting.

For both corpora, we used the named-entity recognizer (NER) in spaCy³ to replace named entities with their NER categories. We replaced tokens outside the top 10,000 most frequent tokens with an “out-of-vocabulary” token.

4.2 Approaches

Throughout our experiments, we compare the following generative models:

1. NEURALEEDITOR: the proposed approach.
2. NLM: a standard left-to-right neural language model generating from scratch. For fair comparison, we use the exact same architecture as the decoder of NEURALEEDITOR.
3. KN5: a standard 5-gram Kneser-Ney language model in KenLM (Heafield et al., 2013).
4. MEMORIZATION: generates by sampling a sentence from the training set.
5. SVAE: the sentence variational autoencoder of Bowman (Bowman et al., 2016), sharing

³honnibal.github.io/spaCy

the same decoder architecture as NEURALEEDITOR. We compare the edit vector to differences of sentence vectors in the SVAE.⁴

4.3 Generative modeling

Perplexity. We start by evaluating NEURALEEDITOR’s value as a language model, measured in terms of perplexity. For the NEURALEEDITOR, we use the likelihood lower bound in Equation 3 where we sum over training set instances within Jaccard distance < 0.5 , and for the VAE term in NEURALEEDITOR, we use the one-sample approximation to the lower bound used in Kingma (2014) and Bowman (2016).

Compared to NLM, our initial result is that NEURALEEDITOR is able to drastically improve log-likelihood for a significant number of sentences in the test set (Figure 2) when considering test sentences with at least one similar sentence in the training set. However, it places lower log-likelihood and on sentences which are far from any prototypes, as it was not trained to make extremely large edits. Proximity to a prototype seems to be the chief determiner of NEURALEEDITOR’s performance.

To evaluate NEURALEEDITOR’s perplexity, we use smoothing with NLM to account for rare sentences not within our Jaccard distance threshold.⁵ We find NEURALEEDITOR improves perplexity over NLM and KN5. Table 1 shows that this is the case for both YELP and the more general BILLIONWORD, which contains substantially fewer test-set sentences close to the training set. On YELP, we surpass even the best ensemble of NLM and KN5, while on BILLIONWORD we nearly match their performance.

Since NEURALEEDITOR draws its strength from sentences in the training set, we also compared against a simpler alternative, in which we ensemble the NLM and retrieval without edits (MEMORIZATION). NEURALEEDITOR performs dramatically better than this alternative. Table 2 also qualitatively demonstrates that sentences generated by NEURALEEDITOR are substantially different from the original prototypes.

⁴We followed the procedure of Bowman (2016) with the modification that the KL penalty weight is annealed to 0.9 instead of 1.0 to avoid encoder degeneracy.

⁵We smooth with NLM, as this corresponds to the model which does not condition on a prototype at all. Smoothing was done with linear interpolation, and weights were selected via a held-out validation set.




Figure 2: The NEURALEEDITOR outperforms the NLM on examples similar to those in the training set (left panel, point size indicates number of training set examples with Jaccard distance < 0.5). The N-gram baseline (right) shows no such behavior, with the NLM outperforming KN5 on most examples.

Model	Perplexity (Yelp)	Perplexity (BILLIONWORD)
KN5	56.546	78.361
KN5+MEMORIZATION	55.180	73.468
NLM	40.174	55.146
NLM+MEMORIZATION	38.980	50.969
NLM+KN5	38.149	47.472
NEURALEEDITOR($\kappa = 0$)	27.600	48.755
NEURALEEDITOR($\kappa = 25$)	27.480	48.921

Table 1: Perplexity of the NEURALEEDITOR with the two VAE parameters κ outperform all methods on YELP and all non-ensemble methods on BILLIONWORD.

Human evaluation. We now turn to human evaluation of generation quality, focusing on grammaticality and plausibility.⁶ We evaluate generations from NEURALEEDITOR against a NLM with a temperature parameter on the per-token softmax⁷, which is a popular technique for suppressing incoherent and ungrammatical sentences. Many NLM systems have noted a undesirable tradeoff between grammaticality and diversity, where a temperature low enough to enforce grammaticality results in short and generic utterances (Li et al., 2016).

Figure 3 illustrates that both the grammaticality and plausibility of NEURALEEDITOR with a temperature of 1.0 is on par with the best tuned temperature for NLM, with a far higher diversity, as measured by unigram entropy. We also find that decreasing the temperature of the NEURALEEDITOR can be used to slightly improve the grammaticality, without substantially reducing the diversity of the generations.

A key advantage of edit-based models thus emerges: Prototypes sampled from the training set

⁶Human raters were asked, “How plausible is it for this sentence to appear in the corpus?” on a scale of 1–3.

⁷If s_i is the softmax logit for token w_i and τ is a temperature parameter, the temperature-adjusted distribution is $p(w_i) \propto \exp(s_i/\tau)$.

Prototype x'	Revision x
i had the fried whitefish taco which was decent, but i've had much better.	i had the <unk> and the fried carnitas tacos, it was pretty tasty, but i've had better.
"hash browns" are unseasoned, frozen potato shreds burnt to a crisp on the outside and mushy on the inside.	the hash browns were crispy on the outside, but still the taste was missing.
i'm not sure what is preventing me from giving it <cardinal> stars, but i probably should.	i'm currently giving <cardinal> stars for the service alone.
quick place to grab light and tasty teriyaki.	this place is good and a quick place to grab a tasty sandwich.
sad part is we've been there before and its been good.	i've been here several times and always have a good time.

Table 2: Edited generations are substantially different from the sampled prototypes.

organically inject diversity into the generation process, even if the temperature of the decoder in the NEURALEEDITOR is zero. Hence, we can keep the decoder at a very low temperature to maximize grammaticality and plausibility, without sacrificing sample diversity. In contrast, a zero temperature NLM would collapse to outputting one generic sentence.

This also suggests that the temperature parameter for the NEURALEEDITOR captures a more natural notion of diversity — higher temperature encourages more aggressive extrapolation from the training set while lower temperatures favor more conservative mimicking. This is likely to be more useful than the tradeoff for generation-from-scratch, where the temperature also affects the quality of generations.

4.4 Semantics of the neural editor

In this section, we investigate learned semantics of the NEURALEEDITOR, focusing on the two desiderata discussed in Section 2: semantic smoothness,




Figure 3: NEURALEEDITOR provides plausibility and grammaticality on par with the best, temperature-tuned language model without any loss of diversity as a function of temperature. Results are based on 400 human evaluations.

and consistent edit behavior.

In order to establish a baseline for these properties, we consider existing sentence generation techniques which can sample semantically similar sentences. We are not aware of other approaches which attempt to learn a vector space for edits, but there are many approaches which learn a vector space for sentences. Of particular relevance is the sentence variational autoencoder (SVAE) which also imposes semantic structure onto a latent vector space, but uses the latent vector to represent the entire sentence, rather than just an edit. We use the SVAE as a baseline by using perturbations in the latent sentence vector and decoding as a way to generate similar/“edited” sentences.

Semantic smoothness. A good editing system should have fine-grained control over the semantics of a sentence: i.e., each edit should only alter the semantics of a sentence by a small and well-controlled amount. We call this property semantic smoothness.

To study smoothness, we first generate an “edit sequence” by randomly selecting a prototype sentence, and then repeatedly editing via the neural editor (with edits drawn from the edit prior $p(z)$) to produce a sequence of revisions. We then ask human annotators to rate the size of the semantic changes between revisions. An example is given in Table 3.

For the SVAE baseline, we generate a similar sequence of sentences by first encoding the prototype sentence, and then decoding after the addition of a random Gaussian with variance 0.4.⁸ This process is repeated to produce a sequence of sentences which we can view as the SVAE equivalent of the edit sequence.

⁸The variance was selected so that the SVAE and NEURALEEDITOR have the same average human similarity judgement between two successive sentences. This avoids situations where the SVAE produces completely unrelated sentence due to the perturbation size.




Figure 4: The neural editor frequently generates paraphrases and similar sentences while avoiding unrelated and degenerate ones. In contrast, the SVAE frequently generates identical and unrelated sentences and rarely generates paraphrases.¹⁰

Figure 4 shows that the neural editor frequently generates paraphrases despite being trained on lexical similarity, and only 1% of edits are unrelated from the prototype. In contrast, the SVAE often repeats sentences exactly, and when it makes an edit it is equally likely to generate unrelated sentences. This suggests that the neural editor produces substantially smoother sentence sequences with a surprisingly high frequency of paraphrases.

Qualitatively (Table 3), NEURALEEDITOR seems to generate long, diverse sentences which smoothly change over time, while the SVAE biases towards short sentences with several semantic jumps, presumably due to the difficulty of training a sufficiently informative SVAE encoder.

Smoothly controlling sentences. We now show that we can selectively choose edits sampled from the neural editor to incrementally optimize a sentence towards desired attributes. This task serves as a useful measure of semantic coverage: if an edit model has high coverage of sentences that are se-

¹⁰545 similarity assessments pairs were collected through Amazon Mechanical Turk following Agirre (2014), with the same scale and prompt. Similarity judgements were converted to descriptions by defining Paraphrase (5), Roughly Equivalent (4-3), Same Topic (2-1), Unrelated (0).

NEURALEEDITOR	SVAE
this food was amazing one of the best i've tried, service was fast and great.	this food was amazing one of the best i've tried, service was fast and great.
this is the best food and the best service i've tried in <gpe>.	this place is a great place to go if you want a quick bite.
some of the best <norp> food i've had in <date> i've lived in <gpe>.	the food was good, but the service was terrible.
i have to say this is the best <norp> food i've had in <gpe>.	this is the best <norp> food in <gpe>.
best <norp> food i've had since moving to <gpe> <date>.	this place is a great place to go if you want to eat.
this was some of the best <norp> food i've had in the <gpe>.	this is the best <norp> food in <gpe>.
i've lived in <gpe> for <date> and every time we come in this is great	the food was good, the service was great.
i've lived in <gpe> for <date> and have enjoyed my haircut at <gpe> since <date>.	the food was good, but the service was terrible.

Table 3: Example random walks from NEURALEEDITOR, where the top sentence is the prototype.




Figure 5: Neural editors can shorten sentences (left), include common words (center, the word ‘service’) and rarer words (right ‘pizza’) while maintaining similarity.

NEURALEEDITOR	SVAE
the coffee ice cream was one of the best i've ever tried.	the coffee ice cream was one of the best i've ever tried.
some of the best ice cream we've ever had!	the <unk> was very good and the food was good.
just had the best ice - cream i've ever had!	the food was good, but not great.
some of the best pizza i've ever tasted!	the food was good, but not great.
that was some of the best pizza i've had in the area.	the food was good, but the service was n't bad.

Table 4: Examples of word inclusion trajectories for ‘pizza’. The NeuralEditor produces smooth chains that lead to word inclusion, but the SVAE gets stuck on generic sentences.

manically similar to a prototype, it should be able to satisfy the target attribute while deviating minimally from the prototype’s original meaning.

We focus on controlling two simple attributes: compressing a sentence to below a desired length (e.g. 7 words), and inserting a target keyword into the sentence (e.g. “service” or “pizza”).

Given a prototype sentence, we try to discover a semantically similar sentence satisfying the target attribute using the following procedure: First, we generate 1000 edit sequences using the procedure described earlier. Then, we select the sequence with highest likelihood whose endpoint possesses the target attribute. We repeat this process for a large number of prototypes.

We use almost the same procedure for the SVAE, but instead of selecting by highest likelihood, we select the sequence whose endpoint has shortest latent vector distance from the prototype (as this is the

SVAE’s metric of semantic similarity).

In Figure 5, we then aggregate the sentences from the collected edit sequences, and plot their *semantic similarity to the prototype* against their *success in satisfying the target attribute*. Not surprisingly, as target attribute satisfaction rises, semantic similarity drops. However, we also see that the neural editor sacrifices less semantic similarity to achieve the same level of attribute satisfaction as the SVAE. The SVAE is reasonable on tasks involving common words (such as the word *service*), but fails when the model is asked to generate rarer words such as *pizza*. Examples from these word inclusion problems show that the SVAE often becomes stuck generating short, generic sentences (Table 4).

Consistent edit behavior: sentence analogies. In the previous results, we showed that edit models learn to generate semantically similar sentences. We now assess whether the edit vector possesses glob-

Google			Microsoft						Method
gram4-superlative	gram3-comparative	family	JJR_JJS	VB_VBD	VBD_VBZ	NN_NNS	JJ_JJS	JJ_JJR	
0.60	0.92	0.60	0.755	0.635	0.825	0.880	0.02	0.81	GloVe
0.82	0.76	0.36	0.440	0.505	0.570	0.450	0.22	0.41	Edit vector (top 10)
0.07	0.14	0.05	0.260	0.090	0.155	0.130	0.13	0.00	Edit vector (top 1)
0.19	0.07	0.21	0.165	0.190	0.155	0.215	0.10	0.19	Sampling (top 10)

Table 5: Edit vectors capture one-word sentence analogies with performance close to lexical analogies.

	Example 1	Example 2
Context	i've had better service at <org >.	my daughter actually looks forward to going to the dentist!
Edit	+ worst - worse	+ his - her
Result	= best service i've had at <org >.	= my son actually looks forward to going to the dentist!

Table 6: Examples of lexical analogies correctly answered by NEURALEEDITOR. Sentence pairs generating the analogy relationship are shortened to only their lexical differences.

ally consistent semantics. Specifically, applying the same edit vector to different sentences should result in semantically analogous edits.

Formally, suppose we have two sentences, x_1 and x_2 , which are related by some underlying semantic relation r . Given a new sentence y_1 , we would like to find a y_2 such that the same relation r holds between y_1 and y_2 .

Our approach is to estimate the edit vector between x_1 and x_2 as $\hat{z} = f(x_1, x_2)$ — the mode of our edit posterior q . We then apply this edit vector to y_1 using the neural editor to yield $\hat{y}_2 = \operatorname{argmax}_x p_{\text{edit}}(x | y_1, \hat{z})$.

Since it is difficult to output \hat{y}_2 exactly matching y_2 , we take the top k candidate outputs of p_{edit} (using beam search) and evaluate whether the gold y_2 appears among the top k elements.

Evaluation for this task can be difficult, as two arbitrary sentences x_1 and x_2 are often not related by any well-defined relationship. However, prior work already provide well-established sets of word analogies (Mikolov et al., 2013a; Mikolov et al., 2013b). We leverage these to generate a new dataset of sentence analogies, using a simple strategy: given an analogous word pair (w_1, w_2) , we mine the Yelp corpus for sentence pairs (x_1, x_2) such that x_1 is transformed into x_2 by inserting w_1 and removing w_2 (allowing for reordering and inclusion/exclusion of stop words).

For example, given the words (good, best), we mine the sentence pair $x_1 =$ “this was a good restaurant” and $x_2 =$ “this was the best restaurant”. Given

a new sentence $y_1 =$ “The cake was great”, we expect y_2 to be “The cake was the greatest”.

For this task, we initially compared against the SVAE, but it had a top- k accuracy close to zero. Hence, we instead compare to the baseline of randomly sampling an edit vector $\hat{z} \sim p(z)$, instead using \hat{z} derived from $f(x_1, x_2)$. We can also compare to numbers on the original word analogy task, restricted to the relationships we find in the Yelp corpus although these numbers are for the easier task of computing word-level (not sentence) analogies.

Interestingly, the top-10 performance of our model in Table 5 is nearly as good as the performance of GloVe vectors on the simpler lexical analogy task, despite the fact that the sentence prediction task is harder. In some categories the neural editor at top-10 actually performs better than word vectors, since the neural editor has an understanding of which words are likely to appear in the context of a Yelp review. Examples in Table 6 show the model is accurate and captures lexical analogies requiring word reorderings.

5 Related work and discussion

Our work connects with a broad literature on neural retrieval and attention based generation methods, semantically meaningful representations for sentences, and nonparametric statistics.

Neural language models (Bengio et al., 2003) based upon recurrent neural networks and sequence-to-sequence architectures (Sutskever et al., 2014) have been widely used due to their flexibility and

performance across a wide-range of NLP tasks (Jurafsky and Martin, 2000; Kalchbrenner and Blunsom, 2013; Hahn and Mani, 2000; Ritter et al., 2011). Our work is motivated by an emerging consensus that attention based mechanisms (Bahdanau et al., 2015) can substantially improve performance of various sequence to sequence tasks by capturing more information from the input sequence (Wu et al., 2016; Vaswani et al., 2017). Our work extends the applicability of attention mechanisms beyond sequence to sequence tasks by deriving a training method for models which attend to randomly sampled sentences.

There is a growing literature on applying retrieval mechanisms to augment neural sequence-to-sequence models. For example, Song (2016) ensemble a retrieval system and an NLM for dialogue, using the NLM to transform the retrieved utterance, and Gu (2017) used an off-the-shelf search engine system to retrieve and condition on training set examples. Both of these approaches rely on a deterministic retrieval mechanism which selects a prototype x' using some input c . In contrast, our work treats the prototype x' as a latent variable, and marginalizes over all possible prototypes — a challenge which motivates our new lexical similarity training method in Section 3.1. Practically, marginalization over x' makes our model attend to training examples based on similarity of output sequences, while retrieval models attend to examples based on similarity of the input sequences.

In terms of generation techniques that capture semantics, the sentence variational autoencoder (SVAE) (Bowman et al., 2016) is closest to our work in that it attempts to impose semantic structure on a latent vector space. However, the SVAE’s latent vector is meant to represent the entire sentence, whereas the neural editor’s latent vector represents an edit. Our results from Section 4.4 suggest that local variation over edits is easier to model than global variation over sentences.

Our use of lexical similarity neighborhoods is comparable to context windows used in word vector training (Mikolov et al., 2013a). Proximity of words within text and lexical similarity both serve as a filter which reveals semantics through distributional statistics of the corpus. More generally, results in manifold learning demonstrate that a weak metric

such as lexical similarity can be used to extract semantic similarity through distributional statistics (Tenenbaum et al., 2000; Hashimoto et al., 2016).

Prototype-then-edit is a semi-parametric approach that remembers the entire training set and uses a neural editor to generalize meaningfully beyond the training set. The training set provides a strong inductive bias — that the corpus can be characterized by prototypes surrounded by semantically similar sentences reachable by edits. Beyond improvements on generation quality as measured by perplexity, the approach also reveals new semantic structures via the edit vector. A natural next step is to apply these ideas in the conditional setting for tasks such as dialogue generation.

Reproducibility. Our code is available at <https://github.com/anonymous>.

6 Appendix

6.1 Construction of the LSH

The LSH maps each sentence to other lexically similar sentences in the corpus, representing a graph over sentences. To speed up corpus construction, we apply breadth-first search (BFS) over the LSH sentence graph started at randomly selected seed sentences. We store the edges encountered during the BFS, and uniformly sample from this set to form the training set. The BFS ensures that every query to the LSH index returns a valid edit pair, at the cost of adding bias to our training set.

6.2 Neural editor architecture

Encoder. The prototype encoder is a 3-layer biLSTM. The inputs to each layer are the concatenation of the forward and backward hidden states of the previous layer, with the exception of the first layer, which takes word vectors as input. Word vectors were initialized with GloVe (Pennington et al., 2014).

Decoder. The decoder is a 3-layer LSTM with attention. At each time step, the hidden state of the top layer is used to compute attention over the top-layer hidden states of the prototype encoder. The resulting attention context vector is then concatenated with the decoder’s top-layer hidden state and used to compute a softmax distribution over output tokens.

References

- E. Agirre, C. Banea, C. Cardie, D. M. Cer, M. T. Diab, A. Gonzalez-Agirre, W. Guo, R. Mihalcea, G. Rigau, and J. Wiebe. 2014. Semeval-2014 task 10: Multilingual semantic textual similarity. In *International Conference on Computational Linguistics (COLING)*, pages 81–91.
- D. Bahdanau, K. Cho, and Y. Bengio. 2015. Neural machine translation by jointly learning to align and translate. In *International Conference on Learning Representations (ICLR)*.
- Y. Bengio, R. Ducharme, P. Vincent, and C. Jauvin. 2003. A neural probabilistic language model. *Journal of machine learning research*, 3(0):1137–1155.
- S. R. Bowman, L. Vilnis, O. Vinyals, A. M. Dai, R. Jozefowicz, and S. Bengio. 2016. Generating sentences from a continuous space. In *Computational Natural Language Learning (CoNLL)*.
- C. Chelba, T. Mikolov, M. Schuster, Q. Ge, T. Brants, P. Koehn, and T. Robinson. 2013. One billion word benchmark for measuring progress in statistical language modeling. *arXiv preprint arXiv:1312.3005*.
- J. Gu, Y. Wang, K. Cho, and V. O. Li. 2017. Search engine guided non-parametric neural machine translation. *arXiv preprint arXiv:1705.07267*.
- U. Hahn and I. Mani. 2000. The challenges of automatic summarization. *Computer*, 33.
- T. B. Hashimoto, D. Alvarez-Melis, and T. S. Jaakkola. 2016. Word embeddings as metric recovery in semantic spaces. *Transactions of the Association for Computational Linguistics (TACL)*, 4:273–286.
- J. R. Hayes and L. S. Flower. 1986. Writing research and the writer. *American psychologist*, 41(10):1106–1113.
- K. Heafield, I. Pouzyrevsky, J. H. Clark, and P. Koehn. 2013. Scalable modified kneser-ney language model estimation. In *Association for Computational Linguistics (ACL)*, pages 690–696.
- D. Jurafsky and J. H. Martin. 2000. *Speech and language processing: An introduction to natural language processing, computational linguistics, and speech recognition*. Prentice Hall Prentice Hall.
- N. Kalchbrenner and P. Blunsom. 2013. Recurrent continuous translation models. In *Empirical Methods in Natural Language Processing (EMNLP)*.
- D. P. Kingma and M. Welling. 2014. Auto-encoding variational Bayes. *arXiv*.
- P. Koehn, H. Hoang, A. Birch, C. Callison-Burch, M. Federico, N. Bertoldi, B. Cowan, W. Shen, C. Moran, R. Zens, et al. 2007. Moses: Open source toolkit for statistical machine translation. In *Association for Computational Linguistics (ACL)*, pages 177–180.
- J. Li, M. Galley, C. Brockett, J. Gao, and W. B. Dolan. 2016. A diversity-promoting objective function for neural conversation models. In *Human Language Technology and North American Association for Computational Linguistics (HLT/NAACL)*.
- T. Mikolov, K. Chen, G. Corrado, and Jeffrey. 2013a. Efficient estimation of word representations in vector space. *arXiv*.
- T. Mikolov, W. Yih, and G. Zweig. 2013b. Linguistic regularities in continuous space word representations. In *hlt-Naacl*, volume 13, pages 746–751.
- R. Nallapati, B. Zhou, C. Gulcehre, B. Xiang, et al. 2016. Abstractive text summarization using sequence-to-sequence rnns and beyond. *arXiv preprint arXiv:1602.06023*.
- J. Pennington, R. Socher, and C. D. Manning. 2014. Glove: Global vectors for word representation. In *Empirical Methods in Natural Language Processing (EMNLP)*.
- A. Ritter, C. Cherry, and W. B. Dolan. 2011. Data-driven response generation in social media. In *Empirical Methods in Natural Language Processing (EMNLP)*, pages 583–593.
- L. Shao, S. Gouws, D. Britz, A. Goldie, B. Strophe, and R. Kurzweil. 2017. Generating high-quality and informative conversation responses with sequence-to-sequence models. In *Empirical Methods in Natural Language Processing (EMNLP)*.
- Y. Song, R. Yan, X. Li, D. Zhao, and M. Zhang. 2016. Two are better than one: An ensemble of retrieval- and generation-based dialog systems. *arXiv preprint arXiv:1610.07149*.
- I. Sutskever, O. Vinyals, and Q. V. Le. 2014. Sequence to sequence learning with neural networks. In *Advances in Neural Information Processing Systems (NIPS)*, pages 3104–3112.
- J. B. Tenenbaum, V. D. Silva, and J. C. Langford. 2000. A global geometric framework for nonlinear dimensionality reduction. *Science*, pages 2319–2323.
- A. Vaswani, N. Shazeer, N. Parmar, J. Uszkoreit, L. Jones, A. N. Gomez, L. Kaiser, and I. Polosukhin. 2017. Attention is all you need. *arXiv preprint arXiv:1706.03762*.
- A. T. Wood. 1994. Simulation of the von mises fisher distribution. *Communications in statistics-simulation and computation*, pages 157–164.
- Y. Wu, M. Schuster, Z. Chen, Q. V. Le, M. Norouzi, W. Macherey, M. Krikun, Y. Cao, Q. Gao, K. Macherey, et al. 2016. Google’s neural machine translation system: Bridging the gap between human and machine translation. *arXiv preprint arXiv:1609.08144*.
- Yelp. 2017. Yelp dataset challenge. https://www.yelp.com/dataset_challenge.