

Grounded Compositional Outputs for Adaptive Language Modeling

Nikolaos Pappas[♣] Phoebe Mulcaire[♣] Noah A. Smith^{♣◇}

[♣]Paul G. Allen School of Computer Science & Engineering, University of Washington

[◇]Allen Institute for Artificial Intelligence

{npappas, pmulcaire, nasmith@cs.washington.edu}

Abstract

Language models have emerged as a central component across NLP, and a great deal of progress depends on the ability to cheaply adapt them (e.g., through finetuning) to new domains and tasks. A language model’s *vocabulary*—typically selected before training and permanently fixed later—affects its size and is part of what makes it resistant to such adaptation. Prior work has used compositional input embeddings based on surface forms to ameliorate this issue. In this work, we go one step beyond and propose a fully compositional output embedding layer for language models, which is further grounded in information from a structured lexicon (WordNet), namely semantically related words and free-text definitions. To our knowledge, the result is the first word-level language model with a size that does not depend on the training vocabulary. We evaluate the model on conventional language modeling as well as challenging cross-domain settings with an open vocabulary, finding that it matches or outperforms previous state-of-the-art output embedding methods and adaptation approaches. Our analysis attributes the improvements to sample efficiency: our model is more accurate for low-frequency words.

1 Introduction

Language models (LMs) are at the heart of natural language processing, especially following their recent success in the pretraining paradigm (Dai and Le, 2015; Peters et al., 2018; Devlin et al., 2019; Radford et al., 2019, *inter alia*). Continued advances in NLP rely on the adaptability of LMs to domains beyond their training data and to new domains and tasks, e.g., through domain adaptive pretraining followed by finetuning (Gururangan et al., 2020). Here, we focus on an important component of LMs, namely the **output vocabulary**—over which a LM’s probability distribution over the

“next word” (given the history) ranges—and investigate the impact of the type of its representation on the adaptability of neural LMs.

Today, LMs are typically trained with a closed output vocabulary derived from the training data; the vocabulary is not modified when the language model is adapted or deployed. This makes large pre-trained language models struggle with rare words, despite being able to produce contextualized representations for them (Schick and Schütze, 2020). More importantly, this means a generative LM can never give nonzero probability to a specific word it did not see in training. This is a longstanding challenge of language modeling (Jelinek, 1997), but it becomes especially important when we adapt to new domains and tasks.

One way to “open up” the vocabulary is to model sequences of bytes, characters, or “word-pieces” rather than the conventional word tokens (Sennrich et al., 2016; Radford et al., 2018; Ponti et al., 2019). While effective, this approach requires the LM to memorize subsequences if it is to treat them as words. These models appear to require greater network depth and show slower convergence than word-based alternatives (Cherry et al., 2018; Al-Rfou et al., 2019); the extra work comes at a cost. This is one of the reasons why the area of word-level language modeling is still very active (Baevski and Auli, 2019; Sukhbaatar et al., 2019; Khandelwal et al., 2020; Press et al., 2020).

Interpolations between word- and character- or morphology-based LMs represent another class of solutions (Mielke and Eisner, 2018; Gerz et al., 2018; Ataman et al., 2020). These “hybrid” approaches combine benefits from both model types. However, they introduce complexity which makes them potentially more difficult to train, maintain, and analyze. Notable for enabling adaptability are interpolated LMs based on copy mechanisms (Merity et al., 2017), dynamic evaluation (Krause et al.,

2018), and neural caches (Grave et al., 2017c,a); the last provides state-of-the-art adaptation performance and, unlike the rest, it does not require additional training.

We propose a new word-level Grounded Compositional output LM (GroC) that applies a compositional representation to the output vocabulary (Section 3). Each word’s output embedding is built from its surface character sequence and (if available) those of semantically related words and a free-text definition of from WordNet (Fellbaum, 1998). This parameterization offers two chief advantages. First, GroC can assign probability to words not seen during training. This means that a vocabulary different from the training vocabulary—e.g., one associated with a different text domain, crucial in adaptive settings—can be considered at inference time. Second, because there are no word type-specific parameters, the number of model parameters in GroC does not depend on the training vocabulary or its size.

We evaluate GroC on language modeling with both fixed and open vocabularies in English. On standard language modeling (Section 4) we observe that our model has superior perplexity and is more sample efficient than a variety of existing output embedding approaches, including the recent adaptive embedding of Baevski and Auli (2019). The open-vocabulary settings include a cross-domain setting and finetuning (Section 5). We find that GroC also outperforms strong interpolated baselines, including the unbounded neural cache model of Grave et al. (2017a) on “near” domains and performs competitively on “far” domains.

Our analysis shows that our approach has improved sharing across words in the output vocabulary. We show experimentally that the perplexity gains are strongest for low-frequency words, implying improved **sample efficiency** relative to baselines: compositional output representations allow us to predict words from fewer training examples.

2 Preliminaries on Language Modeling

Language models assign probability to sequences of tokens; the task is usually framed as learning the conditional probability distributions over individual tokens given their histories of tokens to the left (Bahl et al., 1983). Training requires a sequence of T tokens $\mathbf{x} = \langle x_1, \dots, x_T \rangle$, each x_t a member of a preselected vocabulary \mathcal{V} . We let $\mathbf{x}_t \in \{0, 1\}^{|\mathcal{V}|}$ denote the one-hot encoding of x_t . The probability

of the sequence \mathbf{x} is factored using the chain rule of probability:

$$p(\mathbf{x}) = \prod_{t=1}^T p(x_t | x_1, \dots, x_{t-1}). \quad (1)$$

To approximate this joint distribution, researchers have fit parametric families based on relative frequencies (Bahl et al., 1983; Kneser and Ney, 1995; Goodman, 2001) and neural networks (Bengio et al., 2003; Mikolov et al., 2010). Here, we focus on the latter due to their established effectiveness (Merity et al., 2018; Baevski and Auli, 2019). Tokens in this work correspond to words but they can also correspond to individual characters (Al-Rfou et al., 2019) or byte pairs (Radford et al., 2019).

2.1 Neural Language Models

To make clear this paper’s contributions, we describe neural language models by decomposing them into several abstract parts.

In most neural language models, the first layer of computation obtains an input embedding of each history word x_j using a lookup function. In our notation, this corresponds to selecting the word type’s row in a fixed **input embedding matrix**, \mathbf{E}^{in} : $\mathbf{x}_j^\top \mathbf{E}^{in}$, which we denote $\mathbf{e}_{x_j}^{in}$. Importantly, however, input embeddings need not be lookups; for example, they can be built compositionally from the characters in the surface form of the word (Ling et al., 2015), an idea central to this work.

Next, the history or “prefix” words $\mathbf{x}_{<t} = \langle x_1, \dots, x_{t-1} \rangle$ is encoded into a fixed, d -dimensional vector \mathbf{h}_{t-1} using a **prefix function** $f : \mathcal{V}^* \rightarrow \mathbb{R}^d$. f can be a recurrent or feedforward network; we will experiment with LSTMs (Hochreiter and Schmidhuber, 1997) in Section 4, but our method is agnostic to the prefix function design. In general, each history encoding is defined as

$$\mathbf{h}_{t-1} = f(\mathbf{e}_{x_1}^{in}, \dots, \mathbf{e}_{x_{t-1}}^{in}). \quad (2)$$

Finally, the distribution over the next word (random variable X_t) is given by

$$p(X_t = x_t | \mathbf{h}_{t-1}) \propto \exp(\mathbf{E}^{out} \mathbf{h}_{t-1} + \mathbf{b}), \quad (3)$$

where $\mathbf{E}^{out} \in \mathbb{R}^{|\mathcal{V}| \times d}$ is the **output embedding matrix** and $\mathbf{b} \in \mathbb{R}^{|\mathcal{V}|}$ is the bias vector (corresponding roughly to unigram log-frequencies of words in the vocabulary).

The parameters of the model—including all parameters of the prefix function f , \mathbf{E}^{in} , \mathbf{E}^{out} , and

\mathbf{b} —are all chosen by maximizing the likelihood of the training sequence \mathbf{x} under the model (Eq. 1). Note that, though we focus on an autoregressive (left-to-right) language model objective, our analysis below is applicable to other language model pretraining objectives such as masked language modeling (Devlin et al., 2019) and replaced token detection (Clark et al., 2020).

2.2 Choice of Output Representations

Above we assumed an output embedding matrix \mathbf{E}^{out} that independently parameterizes each word in the vocabulary with a separate d -dimensional vector. This approach requires $d \times |\mathcal{V}|$ parameters, leading to concerns about cost and overparameterization. Prior work addressed this issue by tying parameters between the input and output embedding matrices (i.e., $\mathbf{E}^{out} = \mathbf{E}^{in}$; Inan et al., 2017; Press and Wolf, 2017). However, the parameters for each word are still independent from each other, as displayed in Figure 1(a).

An alternative, also considered here, is to share output parameters across words as well as with the input embeddings. Specifically, this involves making the output embedding a function of the input embedding using a shared parameterization across words, $\mathbf{E}^{out} = g(\mathbf{E}^{in})$, as displayed in Figure 1(b). For example, Gulordava et al. (2018) used a linear transformation, while Baevski and Auli (2019) used a linear transformation for each frequency bin to dedicate parameters to words proportional to their frequencies. Pappas and Henderson (2019) used a deep residual transformation as g , demonstrating that shared parameterizations perform better than independent ones. The two latter studies also provided evidence that models with shared parameterization are more *sample efficient* than independent parameterizations since they perform better on low-frequency words.

Limitations We argue that dependence of a model’s parameterization on the size of the vocabulary leads to several limitations shared by current word-level language models. First, the output embedding methods above have terms that scale with the vocabulary size, such as the lookup table for the input embedding or the bias vector, which is a concern for the parameterization of infrequent words. Second, handling of words *unseen* in the training data leads us to the convention of uninformative “out-of-vocabulary” word types or linguistically naïve, data-driven vocabulary transformations that

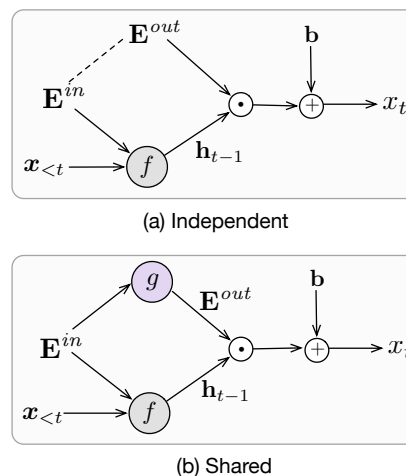


Figure 1: Existing output layer parameterizations.

aggressively decompose words into smaller units (Sennrich et al., 2016). Finally, when pretrained language models are adapted on a downstream task, they do not allow graceful modifications to the vocabulary as required by the task or its data domain. Decoupling the training vocabulary from the target vocabulary that a model can use during inference or finetuning will simplify sequential training and enable open vocabularies.

Building on encouraging results with compositional *input* embeddings (Ling et al., 2015; Józefowicz et al., 2016; Peters et al., 2018), we introduce a language model with shared compositional embeddings for input as well as for *output* word representations. Further, we go beyond past work based on surface forms, making optional use of relations and natural language definitions from structured lexicons like WordNet (Fellbaum, 1998). To our knowledge, this is the first word-level language model whose parameters do not depend on the vocabulary size and which is grounded to an external structured lexicon. Our experiments show that our models are more sample efficient (Section 4) on closed vocabularies and perform competitively on cross-domain settings (Section 5).

3 GroC: Grounded Compositional Output Language Models

We present our grounded compositional output language model (Figure 2).¹ Following the decomposition of neural language models in Section 2 (Equations 2–3), we consider each part of the model in turn: input embeddings (Section 3.1), output embeddings (Section 3.2), and bias (Section 3.3). As noted above, our approach is agnostic to the

¹Code: <https://github.com/Noahs-ARK/groc>

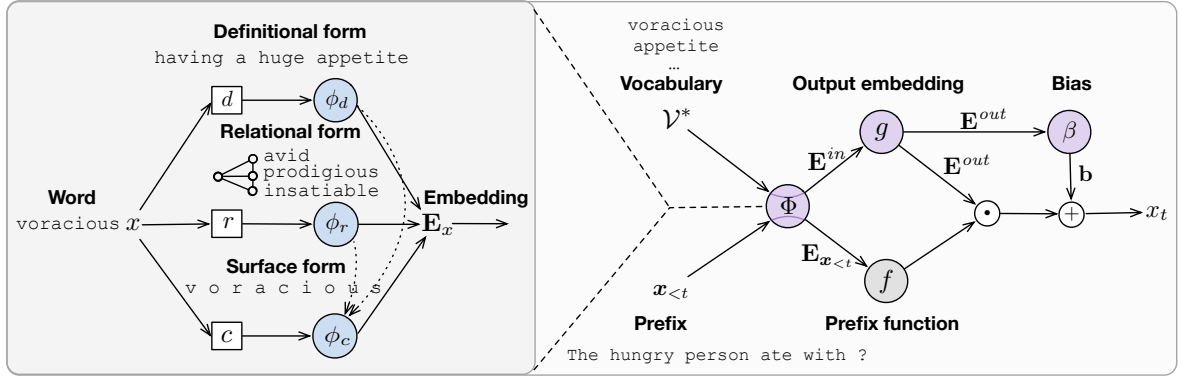


Figure 2: Grounded compositional output language modeling. (Left) The compositional input embedding is grounded in surface, relational, and definitional word forms from an external structured lexicon. (Right) The encoded prefix words are given as input to the prefix function and the words in an arbitrary vocabulary are given as input to the output embedding function and the bias function to predict the next word.

training vocabulary (\mathcal{V}) and to the prefix encoder (f) that has been the focus of most innovations in neural language model design.

3.1 Compositional Input Embeddings

We build on the compositional model of Ling et al. (2015), which encodes a word using its surface string (i.e., character sequence), adding two more sources of information. Peters et al. (2019) enhanced word representations with information from external relational knowledge bases, specifically for words that refer to entities. Like them, we use a structured lexicon (WordNet); we encode every word in the lexicon using its neighbors. The second follows Bahdanau et al. (2017), who used **definitions** to represent out-of-vocabulary words; we encode definitions for all words (regardless of training-set frequency).

We begin by replacing the matrix $\mathbf{E}^{in} \in \mathbb{R}^{|\mathcal{V}| \times d}$ with a neural network that defines a word’s embedding compositionally from its surface form, its position relative to other words in a structured lexicon, and a natural language definition. For each word x , we refer to these, respectively, as the word type’s surface embedding \mathbf{c}_x , relational embedding \mathbf{r}_x , and definitional embedding \mathbf{d}_x . We assume each has a dimensionality of d . The last two are optional (if missing, they are set to zero), and we redefine \mathbf{e}_x as the concatenation of the three, namely $\mathbf{e}_x = \langle \mathbf{c}_x, \mathbf{r}_x, \mathbf{d}_x \rangle$. For \mathbf{r}_x and \mathbf{d}_x , we used the structured relations (synonyms and hyponyms) and free-text definitions in WordNet (Fellbaum, 1998).

In this study, we focus on simple, computationally efficient options for the three encoders. A word x ’s character sequence is encoded as surface encoding \mathbf{c}_x using a convolutional network followed by

a highway network (Józefowicz et al., 2016; Peters et al., 2018). Its relational encoding \mathbf{r}_x is given by an average of $\mathbf{c}_{x'}$ across WordNet synonyms and hyponyms x' . The definitional encoding of x , \mathbf{d}_x , we similarly take an average of the surface encodings $\mathbf{c}_{x'}$ over words x' appearing in the definition. For computational efficiency, we set a maximum limit to the number of words to be used for both relations and definitions (see Appendix B.1).

If a word’s information is not in WordNet, we set \mathbf{r}_x and/or \mathbf{d}_x to $\mathbf{0}$. In future work, additional encodings could be appended, such as contextualized examples (Khandelwal et al., 2020).

A notable property of these input embeddings is that their parameter count does not depend on the vocabulary size $|\mathcal{V}|$. Further, the vocabulary used in training need not be identical to the one used during finetuning, evaluation, or deployment. For example, during training we can use the full vocabulary combined with a softmax approximation method (e.g., Grave et al., 2017b), or by dynamically narrowing the choice of x_t based on its history using co-occurrence statistics (L’Hostis et al., 2016). During finetuning or evaluation, one can use the same vocabulary (required for traditional perplexity evaluations) or a different one chosen statically or dynamically, since any word’s input embedding can be calculated compositionally.

3.2 Compositional Output Embeddings

One straightforward option for vocabulary size-independent output embeddings is to reuse the compositional input embeddings from Section 3.1, along the lines of Press and Wolf (2017). Concretely, at timestep t , we take the set \mathcal{V}'_t of output word types allowed, embed each word type $v \in \mathcal{V}'_t$

as in Section 3.1, and stack these into a matrix \mathbf{E}_t^{in} which serves directly as \mathbf{E}_t^{out} .

Though these compositional representations do enable extensive sharing across the vocabulary, we suspect that the features they capture may require additional processing before capturing “output” distributional similarity, especially when another domain is the real target use case for the language model. This follows prior work discussed in Section 2.2, which showed that making the output embedding a function of the input embeddings with shared parameters improves over simple tying.²

We therefore adopt a depth- k residual network for the output embedding function g (from Section 2.2) that consists of a feedforward function g_i at each layer j with d -dimensions each and apply it to the input embedding at timestep t :

$$\forall j : 1 \leq j \leq k, \mathbf{E}_t^{out(j)} = g_j \left(\mathbf{E}_t^{out(j-1)} \right) + \mathbf{E}_t^{in} \\ \mathbf{E}_t^{out(0)} = \mathbf{E}_t^{in}. \quad (4)$$

Hence, we use $\mathbf{E}_t^{out(k)}$ as the output embedding at timestep t . To avoid overfitting, we apply variational dropout in between the layers, following Pappas and Henderson (2019). In contrast to that work, our resulting output embeddings are compositional. The depth k and the dropout rate are hyperparameters to be tuned on development data. The number of parameters is proportional to k times the number of parameters in the feedforward network ($O(d^2)$); it does not depend on the vocabulary size.

3.3 Bias

In conventional language models, each word in the vocabulary is assigned a bias parameter that roughly captures its log-frequency under a unigram distribution. This is the last part of a neural language model whose parameters depend on the vocabulary size. Instead of a dedicated, independent bias parameter for each word $v \in \mathcal{V}$, we define

$$b_v = \sigma \left(\mathbf{w} \cdot \mathbf{e}_v^{out} + a \right), \quad (5)$$

where σ is the activation function and we introduce parameters $\mathbf{w} \in \mathbb{R}^d$ and $a \in \mathbb{R}$. The bias values b_v are stacked to form \mathbf{b} and used in Equation 3.

3.4 Training

Since all components are differentiable with respect to their parameters, the entire model can be trained

²Note that the input embeddings are passed through the prefix encoder f , which uses additional parameters to create the hidden state \mathbf{h}_{t-1} .

to maximize training-data likelihood as described earlier (Section 2.1). Parameters include:

- Input character embeddings, the convolutional network for \mathbf{c}_* , and $3d^2$ parameters for projection (Section 3.1);
- Output embedding transformation, including the depth- k feedforward network for output embeddings (Section 3.2) and the bias parameters (Section 3.3); and
- Prefix encoder f , an orthogonal design choice to our method (an LSTM in our experiments).

The model size can be adjusted by changing output embedding hyperparameters to fit a given memory requirement — this is the same as any other neural network. Note that despite our vocabulary-size independent parameterization, we still need to process all the words in the supplied vocabulary leading to increased training times despite the model’s sample efficiency. This can be prohibitive for very large vocabularies ($\geq 100K$), where we recommend using softmax approximation methods and making sparse updates of the output embedding parameters (see Appendix 1.3). During inference, \mathbf{E}^{out} can be cached for fast access; there is no need to execute a forward pass more than once.

4 Conventional Language Modeling

We first establish the performance of GroC in the conventional closed-vocabulary setting, considering two datasets. We consider out-of-sample generalization (measured by test-set perplexity) and also analyze fit across the vocabulary by frequency bin.

4.1 Experimental Setup

Datasets. We evaluate our methods on two English datasets: penn (Marcus et al., 1993) and wikitext2 (Merity et al., 2017). We report test perplexity using the provided training/dev./test splits (see details in Appendix B.3). Table 1 also quantifies the percentage of each dataset’s vocabulary that is covered by WordNet (used to derive relational and definitional encodings).

Dataset	genre	$ \mathcal{V} $	# tokens	WNet cov.	
penn	news	10K	929K	78	86
wikitext2	Wikip.	33K	2M	73	76

Table 1: Language modeling dataset statistics.

Models. All of the models compared use the same prefix encoder: a vanilla recurrent neural network based on the implementation by Merity et al. (2017) with 2 layers and 1024 LSTM units, regularized with hidden unit dropout of 0.65 along the lines of Grave et al. (2017a). More details are given in Appendix B.4. The following output embedding approaches are compared:

- **Lookup table:** trains a full output embedding lookup table that corresponds to the vocabulary as defined in Eq. 3.
- **Convolutional** (Józefowicz et al., 2016): an alternative to a lookup table that uses a character-level convolutional neural network followed by a highway network plus a linear “correction” for each vocabulary element to represent the outputs.³
- **Tied** (Press and Wolf, 2017): avoids training separate input and output embedding matrices by tying their parameters. This is a common technique that mitigates the overparameterization issue of the lookup table.
- **Bilinear** (Gulordava et al., 2018): performs a simple linear transformation of the input embedding to produce the output embedding that effectively shares parameters across outputs.
- **Deep residual** (Pappas and Henderson, 2019): performs a deep residual transformation of the input embedding with variational dropout in between its layers, which is more expressive than the bilinear one.
- **Adaptive** (Baevski and Auli, 2019): uses a bilinear transformation of the input and output embedding with parameters proportional to the word frequencies, to assign more capacity to frequent words and less capacity to infrequent ones. This is considered to be a state-of-the-art output embedding method.
- **GroC** (ours): the grounded compositional output embedding described in Section 3.2.

For fair comparison, we apply variational dropout to all output embeddings. Hyperparameter selection of dropout rates, output network depth and activation, linear “correction,” and adaptive frequency cutoffs was conducted by grid search on validation data. Details are given in Appendix B.2.

³Note that we chose not to use a linear “correction” with GroC since it deviates from our goal of having a vocabulary-independent parameterization, but it could be applied to GroC in the future for additional improvements.

Output embedding	penn		wikitext2	
	$ \Theta $	test	$ \Theta $	test
Lookup table	13M	90.8	23M	108.3
Convolutional [J16]	13M	101.6	23M	116.6
Tied [PW17]	10M	86.2	15M	97.3
Bilinear [G18]	10M	83.7	15M	95.9
Deep residual [PH19]	10M	80.5	15M	94.7
Adaptive [AM19]	8M	79.3	9M	90.7
GroC (ours)	9M	69.5	9M	82.5

Table 2: Perplexity scores on conventional language modeling benchmarks with closed vocabulary. $|\Theta|$ denotes the total number of model parameters.

4.2 Results

Table 2 reports perplexities achieved by all seven models. The main finding is that GroC achieves lower perplexity than the previous models, on both datasets. Note that GroC outperforms the state-of-the-art output embedding method of Baevski and Auli (2019); specifically, by -9.8 and -8.2 points on penn and wikitext-2, respectively. The difference with the other methods is even larger. We also confirm the findings of Pappas and Henderson (2019), that output parameter sharing methods outperform tied output embedding and the lookup table, and, of Józefowicz et al. (2016), that convolutional output embeddings lag behind full softmax (lookup table). Notably, GroC outperforms the best reported scores by Merity et al. (2017) and Grave et al. (2017a) on penn, using about 11M fewer parameters and a similar prefix network to the latter. See Appendix B.5 for a more detailed comparison with state-of-the-art models of similar size.

Nevertheless, GroC is about $1.3\times$ slower than the convolutional method on penn; with sparse updates ($p > 0.3$) we can make it $2.1\times$ faster than that method, which is comparable to the speed of the bilinear method, while maintaining a perplexity improvement of -26 points (see detailed speed comparisons in Table 10 in Appendix B.4).

4.3 Analysis

The experiment above establishes that our approach achieves improved perplexity relative to alternative output embeddings. We next decompose its performance in various ways to understand why.

Word frequency effects. We conjecture that GroC’s main benefit comes from words that are rare in the training data, since the core contribution is to share representations across the vocabulary. To evaluate this hypothesis, we consider the difference in test loss (cross entropy) between GroC and a




Figure 3: Median loss difference between each baseline and GroC over different word frequency intervals on `penn` (a) and `wikitext2` (b). The biggest differences are mostly observed on words with low training frequencies.

baseline model, following Baevski and Auli (2019) but computing the median instead of the average to reduce the effect of outliers. We decompose this score by data frequency bins (e.g., words occurring 1–50 times in the training dataset). Figure 3 is displayed for the `penn` and `wikitext2` datasets. The trend we observe is that GroC has the greatest relative benefit for words in lower frequency bins, compared to each model. The lowest-frequency bin on `penn` deviates from this pattern, which we take as an indication that generalizing to infrequent words with only 1M training tokens and a small 10K vocabulary is inherently challenging.

Ablations. To assess the contributions of GroC’s components, we performed ablation tests on `penn` and `wikitext2` (Table 3). These include removing relational and/or definitional forms, either with or without a deep residual output network. For fairness, we tune the hyperparameters of the ablated model variants as above. Overall, removing the relational and definitional forms from the main model with or without output network on top increases the perplexity. The largest drop in perplexity happens when we remove both forms, which highlights their notable contribution to the full model. Lastly, the results on `wikitext2` highlight the importance of capturing the output similarity with an output network (out) for datasets with a larger vocabulary as opposed to merely reusing the grounded compositional embeddings as output embeddings.

Model	penn		wikitext2	
	dev.	test	dev.	test
GroC + out	75.0	71.4 -	87.0	82.5 -
– relations	77.1	72.7 ↑	90.2	85.3 ↑
– definitions	75.6	72.0 ↑	88.6	84.3 ↑
– both	79.8	75.8 ↑	94.3	89.8 ↑
GroC	72.5	69.5 -	88.7	84.1 -
– relations	74.2	70.8 ↑	93.0	88.0 ↑
– definitions	74.4	71.1 ↑	87.6	83.1 ↓
– both	76.3	73.2 ↑	94.5	89.5 ↑

Table 3: Ablated model variants on `penn` and `wikitext-2`. *out*: the deep residual output network,

Lexicon coverage. To measure the effect of lexicon coverage on model performance in a controlled setting, we artificially remove words from WordNet, making them unavailable for relational and definitional encodings. In this experiment, we consider the `penn` dataset, where WordNet’s coverage over the (relatively small) vocabulary is highest to begin with. Table 4 shows the resulting test perplexity of a pretrained model (inference) and a model trained from scratch (train) when such controlled manipulation is applied to them from 0% up to the maximum of 82% coverage (Table 1). Note that we treat relational independently of definitional forms since they are not always co-present. Overall, the results indicate that the model is sensitive to changes in the forms of words that have been seen during training but it is robust to changes if it is trained from scratch. In the next section, we

Coverage	surf.	0%	16%	32%	48%	64%	82%
inference	73.1	187.8	159.3	128.5	102.6	83.5	69.5
train	–	72.4	70.0	70.4	69.6	70.7	69.5

Table 4: External lexicon coverage effect on the perplexity of GroC on the `penn` test set. *surf.*: model with surface forms only from Table 3, last row.

investigate what happens when we add forms for words which the model has never seen before.

5 Cross-Domain Language Modeling

To demonstrate our model’s ability to generalize beyond its training data, we evaluate it across domains with an open vocabulary, in two settings: zero-resource, where it is first trained on one domain and then tested on a new target domain, and low-resource, in which the model is further exposed to training data in the new domain.

5.1 Experimental Setup

Data. Following Grave et al. (2017a), we create English datasets from News Crawl (Bojar et al.,

Model		near domains					far domains	
		2007 →	2008	2009	2010	2011	Web	Wiki
Zero-resource setting								
Grave et al. (2017a)	Base	220.9	237.6	256.2	259.7	268.8	689.3	1003.2
	Base + unigram	220.3	235.9	252.6	256.1	264.3	581.1	609.4
	Base + local cache	218.9	234.5	250.5	256.2	265.2	593.4	316.5
	Base + unbounded cache	166.5	191.4	202.6	204.8	214.3	383.4	337.4
Baseline	Tied	184.3	199.8	217.3	221.6	229.9	660.6	841.1
	Tied + unigram [G17]	187.8	203.6	221.5	225.9	234.3	577.5	819.7
	Tied + local cache [G17]	181.8	196.5	212.0	217.7	225.9	501.7	406.9
Ours	GroC	158.6	171.0	186.7	192.5	200.4	637.9	753.9
	GroC + unigram [G17]	155.2	167.3	183.1	189.5	196.4	533.6	689.2
	GroC + local cache [G17]	152.6	164.1	179.0	185.1	192.3	493.0	408.8
Low-resource setting								
	Tied + finetuning	–	172.8	177.9	180.7	185.4	212.7	242.6
	GroC + finetuning	–	153.7	162.2	167.0	170.6	239.5	216.9

Table 5: Results on *near* and *far* cross-domain language modeling with an open vocabulary with a zero-resource or a low-resource setting. Top four rows display scores from Grave et al. (2017a), while the next three are from our re-implementation with a stronger base model. Boldface marks the best perplexity on each test set.

Dataset	source	train V	test V	OOV%
2007		81K	188K	2.0
2008		82K	197K	2.3
2009	News Crawl	81K	195K	2.5
2010		78K	181K	2.4
2011		80K	184K	2.5
web	Common Crawl	75K	174K	5.8
wiki	WikiText-103	67K	109K	5.4

Table 6: Dataset statistics for cross-domain experiments. OOV% gives the percentage of tokens in the test set not present in the 2007 train vocabulary.

2014), Common Crawl,⁴ and WikiText-103 (Merity et al., 2017). Dataset statistics are given in Table 6. All models are trained on 2M tokens from the 2007 dataset and evaluated on 10M tokens; finetuning is done on an additional 2M tokens from the target domain. We consider the domain of the 2008–2011 datasets to be similar (“near”) to that of the training set, 2007, as they contain news from different time periods. In comparison, *web* and *wiki* are more different (“far”) from 2007.

Models. We compare GroC to the tied output embedding model described in Section 4.1 when combined with the following adaptation methods:

- **Unigram:** we interpolate the model’s distribution with a unigram cache, which assigns probabilities based on the counts of words in the test data observed so far during evaluation.
- **Local cache:** we interpolate the model’s dis-

tribution with a neural cache (Grave et al., 2017c), which assigns probabilities based on the similarity of the current hidden state to previous hidden states during evaluation.

- **Finetuning:** the model is finetuned on 2M tokens from the target domain.

(We also compare to the reported unbounded cache results from Grave et al., 2017a.) Cache models provide effective adaptation without training by using recent history to develop an auxiliary distribution during evaluation, informing predictions of unseen or rarely-seen words. However, as GroC already assigns non-negligible weight to new words not seen prior to evaluation, the cache has less effect by default, even if its predictions are more accurate, an effect we observed in validation. To address this, we down-weighted the model’s predictions for new words prior to cache interpolation by 0.1. For finetuning, both tied and GroC models were trained for an additional 3 epochs on the target domain, allowing them to adapt to the new domain. See Appendix C.4 for hyperparameter details.

Vocabulary setting. For a fair comparison, all models are evaluated on the union of the training and test vocabularies. Tied models are interpolated with the uniform distribution at test time to prevent infinite perplexities on unseen words, prior to cache interpolation if applicable. Words present in the finetuning data but not in the original training data are given random embeddings prior to finetuning.

⁴We used the version from WMT 2014 (Bojar et al., 2014).

5.2 Results

The results for the cross-domain experiments are shown in Table 5. Standalone GroC improves perplexity relative to the tied model in every domain by up to -30 points, including the local neural cache and the unbounded neural cache model in the near-domain, even when the former is applied to our own stronger tied-embedding baseline model. In addition, finetuned GroC outperforms all previous baselines by a wide margin including the unbounded cache by about -40 and -132 points on near and far domains, respectively. Here, GroC outperforms the finetuned tied model by up to -25 points except in `web` domain, and reaches lower validation scores with fewer iterations in 5 out of 6 domains (see Appendix C.1). For the `web` domain, caches and finetuning are more effective than in any other domain, indicating unique domain dynamics worthy of further study.

6 Conclusion

We proposed an adaptive language model based on grounded compositional outputs. We demonstrated that it reduces the number of parameters and increases sample efficiency, outperforming strong output embedding methods and adaptation baselines on both in-domain and open-vocabulary settings respectively. In principle, our results should be applicable to word-piece language models which are currently based on lookup tables to improve their sample efficiency and compactness. In future work, it would be interesting to investigate to what extent pretrained language models benefit from GroC on such zero-resource or low-resource adaptation settings. This work indicates several other future directions for language modeling in low-resource domains: extension to other languages, scaling training to even larger vocabularies, and applying GroC in a large pretraining setting to expand its zero-shot generalization.

Acknowledgments

We thank Sofia Serrano, Ofir Press, Yizhong Wang, Mourad Heddaya, Kelvin Luu, and the anonymous reviewers for their helpful feedback. Nikolaos Pappas was supported by the Swiss National Science Foundation under the project UNISON, grant number P400P2_183911. This work was supported in part by US NSF grant 1562364 and the Office of Naval Research under MURI grant N00014-18-1-2670.

References

- Rami Al-Rfou, Dokook Choe, Noah Constant, Mandy Guo, and Llion Jones. 2019. [Character-level language modeling with deeper self-attention](#). In *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence*, Honolulu, HI, USA.
- Duygu Ataman, Wilker Aziz, and Alexandra Birch. 2020. [A latent morphology model for open-vocabulary neural machine translation](#). In *Proceedings of the International Conference on Learning Representations*.
- Alexei Baevski and Michael Auli. 2019. [Adaptive input representations for neural language modeling](#). In *Proceedings of the International Conference on Learning Representations*, New Orleans, LA, USA.
- Dzmitry Bahdanau, Tom Bosc, Stanislaw Jastrzebski, Edward Grefenstette, Pascal Vincent, and Yoshua Bengio. 2017. [Learning to compute word embeddings on the fly](#). *CoRR*, abs/1706.00286.
- L. R. Bahl, F. Jelinek, and R. L. Mercer. 1983. [A maximum likelihood approach to continuous speech recognition](#). *IEEE Transactions on Pattern Analysis and Machine Intelligence*, PAMI-5(2):179–190.
- Yoshua Bengio, Réjean Ducharme, Pascal Vincent, and Christian Janvin. 2003. [A neural probabilistic language model](#). *Journal of Machine Learning Research*, 3:1137–1155.
- Ondřej Bojar, Christian Buck, Christian Federmann, Barry Haddow, Philipp Koehn, Johannes Leveling, Christof Monz, Pavel Pecina, Matt Post, Herve Saint-Amand, et al. 2014. [Findings of the 2014 Workshop on Statistical Machine Translation](#). In *Proceedings of the Ninth Workshop on Statistical Machine Translation*, pages 12–58.
- Colin Cherry, George Foster, Ankur Bapna, Orhan Firat, and Wolfgang Macherey. 2018. [Revisiting character-based neural machine translation with capacity and compression](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 4295–4305, Brussels, Belgium. Association for Computational Linguistics.
- Kevin Clark, Minh-Thang Luong, Quoc V. Le, and Christopher D. Manning. 2020. [ELECTRA: Pre-training text encoders as discriminators rather than generators](#). In *Proceedings of the International Conference on Learning Representations*.
- Andrew M Dai and Quoc V Le. 2015. [Semi-supervised sequence learning](#). In C. Cortes, N. D. Lawrence, D. D. Lee, M. Sugiyama, and R. Garnett, editors, *Advances in Neural Information Processing Systems 28*, pages 3079–3087. Curran Associates, Inc.
- Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. 2019. [BERT: Pre-training of deep bidirectional transformers for language understanding](#). In *Proceedings of the 2019 Conference*

- of the North American Chapter of the Association for Computational Linguistics: *Human Language Technologies, Volume 1 (Long and Short Papers)*, pages 4171–4186, Minneapolis, Minnesota. Association for Computational Linguistics.
- Christiane Fellbaum, editor. 1998. *WordNet: An Electronic Lexical Database*. MIT Press.
- Yarin Gal and Zoubin Ghahramani. 2016. [A theoretically grounded application of dropout in recurrent neural networks](#). In *Proceedings of the 30th International Conference on Neural Information Processing Systems*, pages 1027–1035, USA. Curran Associates Inc.
- Daniela Gerz, Ivan Vulić, Edoardo Ponti, Jason Naradowsky, Roi Reichart, and Anna Korhonen. 2018. [Language modeling for morphologically rich languages: Character-aware modeling for word-level prediction](#). *Transactions of the Association for Computational Linguistics*, 6:451–465.
- Joshua T. Goodman. 2001. [A bit of progress in language modeling](#). *Computer Speech and Language*, 15(4):403–434.
- Edouard Grave, Moustapha Cisse, and Armand Joulin. 2017a. [Unbounded cache model for online language modeling with open vocabulary](#). In *Proceedings of the 31st International Conference on Neural Information Processing Systems*, pages 6044–6054, USA. Curran Associates Inc.
- Édouard Grave, Armand Joulin, Moustapha Cissé, David Grangier, and Hervé Jégou. 2017b. [Efficient softmax approximation for GPUs](#). In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 1302–1310, International Convention Centre, Sydney, Australia. PMLR.
- Edouard Grave, Armand Joulin, and Nicolas Usunier. 2017c. [Improving neural language models with a continuous cache](#). In *Proceedings of the International Conference on Learning Representations*, Toulon, France.
- Kristina Gulordava, Laura Aina, and Gemma Boleda. 2018. [How to represent a word and predict it, too: Improving tied architectures for language modelling](#). In *Proceedings of the 2018 Conference on Empirical Methods in Natural Language Processing*, pages 2936–2941, Brussels, Belgium. Association for Computational Linguistics.
- Suchin Gururangan, Ana Marasović, Swabha Swayamdipta, Kyle Lo, Iz Beltagy, Doug Downey, and Noah A. Smith. 2020. [Don’t stop pretraining: Adapt language models to domains and tasks](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 8342–8360, Online. Association for Computational Linguistics.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. [Long short-term memory](#). In *Neural Computation*, volume 9 (8), pages 1735–1780. MIT Press.
- Hakan Inan, Khashayar Khosravi, and Richard Socher. 2017. [Tying word vectors and word classifiers: A loss framework for language modeling](#). In *Proceedings of the 5th International Conference on Learning Representations*, Toulon, France.
- Frederick Jelinek. 1997. *Statistical Methods for Speech Recognition*. MIT press.
- Rafal Józefowicz, Oriol Vinyals, Mike Schuster, Noam Shazeer, and Yonghui Wu. 2016. [Exploring the limits of language modeling](#). *CoRR*, abs/1602.02410.
- Urvashi Khandelwal, Omer Levy, Dan Jurafsky, Luke Zettlemoyer, and Mike Lewis. 2020. [Generalization through memorization: Nearest neighbor language models](#). In *Proceedings of the International Conference on Learning Representations*.
- Yoon Kim, Yacine Jernite, David Sontag, and Alexander M. Rush. 2016. [Character-aware neural language models](#). In *Proceedings of the Thirtieth AAAI Conference on Artificial Intelligence*, AAAI’16, pages 2741–2749. AAAI Press.
- R. Kneser and H. Ney. 1995. [Improved backing-off for m-gram language modeling](#). In *1995 International Conference on Acoustics, Speech, and Signal Processing*, volume 1, pages 181–184 vol.1.
- Ben Krause, Emmanuel Kahembwe, Iain Murray, and Steve Renals. 2018. [Dynamic evaluation of neural sequence models](#). In *Proceedings of the 35th International Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Research*, pages 2766–2775, Stockholm, Sweden. PMLR.
- Gurvan L’Hostis, David Grangier, and Michael Auli. 2016. [Vocabulary selection strategies for neural machine translation](#). *CoRR*, abs/1610.00072.
- Wang Ling, Chris Dyer, Alan W Black, Isabel Trancoso, Ramón Fernández, Silvio Amir, Luís Marujo, and Tiago Luís. 2015. [Finding function in form: Compositional character models for open vocabulary word representation](#). In *Proceedings of the 2015 Conference on Empirical Methods in Natural Language Processing*, pages 1520–1530, Lisbon, Portugal. Association for Computational Linguistics.
- Mitchell P. Marcus, Mary Ann Marcinkiewicz, and Beatrice Santorini. 1993. [Building a large annotated corpus of english: The Penn treebank](#). *Computational Linguistics*, 19(2):313–330.
- Stephen Merity, Nitish Shirish Keskar, and Richard Socher. 2018. [Regularizing and optimizing LSTM language models](#). In *International Conference on Learning Representations*, Vancouver, Canada.

- Stephen Merity, Caiming Xiong, James Bradbury, and Richard Socher. 2017. [Pointer sentinel mixture models](#). In *Proceedings of the 5th International Conference on Learning Representations*, Toulon, France.
- Sabrina J. Mielke and Jason Eisner. 2018. [Spell once, summon anywhere: A two-level open-vocabulary language model](#). In *Proceedings of the Thirty-Third AAAI Conference on Artificial Intelligence*, volume abs/1804.08205, Honolulu, HI, USA.
- Tomas Mikolov, Martin Karafiát, Lukás Burget, Jan Cernocký, and Sanjeev Khudanpur. 2010. [Recurrent neural network based language model](#). In *INTER-SPEECH*, pages 1045–1048. ISCA.
- Tomas Mikolov and Geoffrey Zweig. 2012. [Context dependent recurrent neural network language model](#). In *2012 IEEE Spoken Language Technology Workshop (SLT)*, pages 234–239.
- Nikolaos Pappas and James Henderson. 2019. [Deep residual output layers for neural language generation](#). In *Proceedings of the 36th International Conference on Machine Learning*, volume 97 of *Proceedings of Machine Learning Research*, pages 5000–5011, Long Beach, California, USA. PMLR.
- Matthew Peters, Mark Neumann, Mohit Iyyer, Matt Gardner, Christopher Clark, Kenton Lee, and Luke Zettlemoyer. 2018. [Deep contextualized word representations](#). In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 1 (Long Papers)*, pages 2227–2237, New Orleans, Louisiana. Association for Computational Linguistics.
- Matthew E. Peters, Mark Neumann, Robert Logan, Roy Schwartz, Vidur Joshi, Sameer Singh, and Noah A. Smith. 2019. [Knowledge enhanced contextual word representations](#). 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 43–54, Hong Kong, China. Association for Computational Linguistics.
- Edoardo Maria Ponti, Ivan Vulić, Ryan Cotterell, Roi Reichart, and Anna Korhonen. 2019. [Towards zero-shot language modeling](#). 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 2893–2903, Hong Kong, China. Association for Computational Linguistics.
- Ofir Press, Noah A. Smith, and Omer Levy. 2020. [Improving transformer models by reordering their sublayers](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 2996–3005, Online. Association for Computational Linguistics.
- Ofir Press and Lior Wolf. 2017. [Using the output embedding to improve language models](#). In *Proceedings of the 15th Conference of the European Chapter of the Association for Computational Linguistics: Volume 2, Short Papers*, pages 157–163, Valencia, Spain. Association for Computational Linguistics.
- Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. 2018. [Improving language understanding by generative pre-training](#).
- Alec Radford, Jeff Wu, Rewon Child, David Luan, Dario Amodei, and Ilya Sutskever. 2019. [Language models are unsupervised multitask learners](#).
- Timo Schick and Hinrich Schütze. 2020. [BERTRAM: Improved word embeddings have big impact on contextualized model performance](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, pages 3996–4007, Online. Association for Computational Linguistics.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016. [Neural machine translation of rare words with subword units](#). In *Proceedings of the 54th Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pages 1715–1725, Berlin, Germany.
- Sainbayar Sukhbaatar, Edouard Grave, Piotr Bojanowski, and Armand Joulin. 2019. [Adaptive attention span in transformers](#). In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 331–335, Florence, Italy. Association for Computational Linguistics.
- Wojciech Zaremba, Ilya Sutskever, and Oriol Vinyals. 2014. [Recurrent neural network regularization](#). *CoRR*, abs/1409.2329.
- Julian Georg Zilly, Rupesh Kumar Srivastava, Jan Koutník, and Jürgen Schmidhuber. 2017. [Recurrent highway networks](#). In *Proceedings of the 34th International Conference on Machine Learning*, volume 70 of *Proceedings of Machine Learning Research*, pages 4189–4198, International Convention Centre, Sydney, Australia. PMLR.
- Barret Zoph and Quoc V. Le. 2016. [Neural architecture search with reinforcement learning](#). *CoRR*, abs/1611.01578.

A Supplementary Material for “Grounded Compositional Outputs for Adaptive Language Modeling”

We report here the computer infrastructure, experimental details, including hyperparameter bounds, hyperparameter optimal values, training speed, development scores, for both of the experiments where applicable. We also provide a comparison with state of the art by taking into account the number of model parameters and guide the reader through the replication effort we did to reproduce the neural cache by Grave et al. (2017a).

B Conventional Language Modeling

For the experiments with a closed vocabulary on penn⁵ and wikitext-2,⁶ we used the following computing infrastructure: 5 GeForce RTX 2080 Ti gpu cards. Our codebase is based on Pytorch⁷ and is publicly available on Github.⁸

B.1 Model Configuration

The prefix network used by all output embedding methods is a vanilla recurrent neural network based on the implementation by Merity et al. (2017)⁹ with 2 layers and 1024 LSTM units, regularized with hidden unit dropout of 0.65 along the lines of Grave et al. (2017a). The maximum length of the relational and definitional forms from Wordnet is set to 3 and 10 without search based on our computational budget.¹⁰ The embedding size is set to 300 for penn and 256 for wikitext2. For optimization we use Adam with a learning rate of 0.001, initial weight uniformly sampled in the range $[-0.05, 0.05]$, and a batch size of 20 for penn and wikitext2. We clip the norm of the gradient to 0.1 and unroll the network for 35 steps. The learning rate is multiplied by 0.1 if the development loss does not decrease for 4 consecutive epochs and we perform early stopping if there is no improvement for 8 consecutive epochs.

⁵www.fit.vutbr.cz/~imikolov/rnnlm/simple-examples.tgz

⁶s3.amazonaws.com/research.metamind.io/wikitext/wikitext-2-v1.zip

⁷pytorch.org/get-started

⁸github.com/Noahs-ARK/groc

⁹github.com/salesforce/awd-lstm-lm

¹⁰We expect that a larger budget would generally allow to increase these limits and obtain even better results.

Hyperparameter	abbrev.	range	trials
Output dropout	<i>r</i>	{0, 0.1, ..., 1.0}	10
Linear correction	<i>cor</i>	{32, 64, 128}	3
Adaptive cutoffs	<i>cut</i>	{253, 721, 118, 226, 424, 334}	6
Output net depth	<i>k</i>	{0, 1, 2, 3, 4}	4
Output net activation	<i>act</i>	{ <i>relu</i> , <i>selu</i> , <i>tanh</i> }	3

Table 7: Hyperparameters, range of values, and, number of trials required to search them. Adaptive cutoffs are read as follows: e.g. for 253 the cutoff array contains $[0.2 * n, 0.5 * n, 0.3 * n]$, $n = |\mathcal{V}|$ words per bin.

B.2 Hyperparameter Optimization

For all methods, the hyperparameter selection of output embedding dropout rate (*r*), output network depth (*k*) and activation (*act*), linear “correction”, and adaptive frequency cutoffs was conducted by grid search over specific range of values given in Table 7 on development data. Note that not all the hyperparameters apply to all methods, as can be seen in Table 8 where we report the optimal hyperparameter values for each of the methods. For all the baselines we performed exhaustive grid search on both datasets, but for our method we performed grid search only on penn and searched manually on wikitext-2 by selecting values of hyperparameters that were ranked high based on the grid search on penn to avoid the increased cost that comes with training our method (see speed comparison in Appendix B.4). The total number of trials for all methods including our ablations were 204 and 67 respectively for penn and wikitext-2 respectively. Note that the reduced number of trials is due to not performing exhaustive search for our method and its ablations as explained above. The number of trials per method can be derived by

Method	penn					wikitext2				
	<i>r</i>	<i>cor</i>	<i>cut</i>	<i>k</i>	<i>act</i>	<i>r</i>	<i>cor</i>	<i>cut</i>	<i>k</i>	<i>act</i>
Lookup table	0.1	-	-	-	-	0.2	-	-	-	-
Convolutional	0.1	128	-	-	-	0.1	182	-	-	-
Tied	0.0	-	-	-	-	0.0	-	-	-	-
Bilinear	0.5	-	-	-	-	0.4	-	-	-	-
Deep residual	0.5	-	-	4	selu	0.6	-	-	1	selu
Adaptive	0.3	-	2k7k	-	-	0.2	-	6k21k	-	-
GroC (ours)	0.2	-	-	0	-	0.2	-	-	1	relu
- relations	0.3	-	-	0	-	0.3	-	-	1	selu
- definitions	0.2	-	-	0	-	0.3	-	-	2	relu
- both	0.3	-	-	0	-	0.3	-	-	1	selu

Table 8: Best hyperparameter values per method.

multiplying the non-zero columns per row with the number of trials required for each column.

B.3 Development Scores

Table 9 displays the development scores and number of parameters along with the test perplexities for our model and all the baseline output embedding methods for our main experiment. The development scores for the models of the ablation study and for the base models of the coverage experiment have already been given in Table 3 in the main paper (Section 4.3). Overall, we can observe that in most cases the ranking based on the development scores is indicative of the ranking of the methods according to the test scores.

Method	Θ	penn		wikitext2		
		dev.	test	Θ	dev.	test
Lookup table	13M	93.5	90.8	23M	113.8	108.3
Convolutional	13M	104.0	101.6	23M	121.2	116.6
Tied	10M	88.6	86.2	15M	101.0	97.3
Bilinear	10M	87.0	83.7	15M	101.3	95.9
Deep residual	10M	84.0	80.5	15M	100.1	94.7
Adaptive	8M	84.0	79.3	9M	95.8	90.7
GroC (ours)	9M	72.5	69.5	9M	87.0	82.5

Table 9: Development and test scores on conventional language modeling benchmarks with closed vocabulary. $|\Theta|$ denotes the total number of model parameters.

B.4 Training Speed

Table 10 displays the average training speed per epoch in seconds for each of the methods. This experiment was run on a single, dedicated¹¹ GeForce RTX 2080 Ti. As we mentioned in Section 3.1, even though our model has vocabulary-size independent parameterization it is not independent of the computation that is required to encode the vocabulary. This has a negative impact on the training speed of GroC, making it a bit slower than the Convolutional method, namely $1.3\times$ slower.

To mitigate this problem we recommend training GroC with sparse updates for the output embedding parameters as described in the main paper (Section 3.4). Concretely, at each training iteration with probability p we make a full update and keep the output embedding frozen otherwise. The rest of the network is trained with full updates as before. We can observe that this optimization strategy makes GroC nearly as efficient as the base-

¹¹By dedicated GPU card here we mean that no other processes were using the GPU card when we performed the experiments for each of the methods.

lines with $p = 0.1$ or $p = 0.3$. In particular, it becomes even faster than the convolutional baseline by $2.1\times$. Furthermore, our best model with $p = 0.3$ which is much faster reaches 75.3 perplexity on penn without additional hyperparameter optimization which is still -4 points lower than the second best, adaptive output embedding; tuning the model from scratch should likely lead to even better results. This is quite encouraging because it means that the benefits of our model need not come with a large computational cost. In future work, the training speed could be optimized even further by devising specialized efficient training methods for compositional outputs.

Method	penn	wikitext-2
Lookup table	19.5	59.5
Convolutional	201.2	1301.9
Tied	18.6	53.6
Bilinear	35.0	120.1
Deep residual	61.2	114.5
Adaptive	27.2	77.6
GroC (ours)	259.8	1813.5
– 10% updates	236.3	1627.7
– 30% updates	173.5	1262.9
– 50% updates	131.5	936.4
– 70% updates	95.2	669.0
– 90% updates	46.0	299.0

Table 10: Training speed for each method. We report the average time in seconds to complete one epoch.

B.5 Comparison with State-of-the-Art Models

Table 11 displays several state-of-the-art models which have number of parameters ranging from 9M to 20M on Penn Treebank. We can observe that our model which has only 9.7M parameters achieves better performance than all the models that have lower than or equal to 21M parameters and even the model by Inan et al. (2017) which has 24M parameters. Note that our model has lower perplexity than the pointer sentinel mixture model by Merity et al. (2017) and the neural cache model by Grave et al. (2017a) while having 11M less parameters than them.

Moreover, it is very close to the other models which have around 23–25M parameters without being highly regularized (weight dropout, input dropout) or having advanced optimization strategies (SGD + ASGD, finetuning) like AWD-LSTM (Merity et al., 2017). Training larger models and

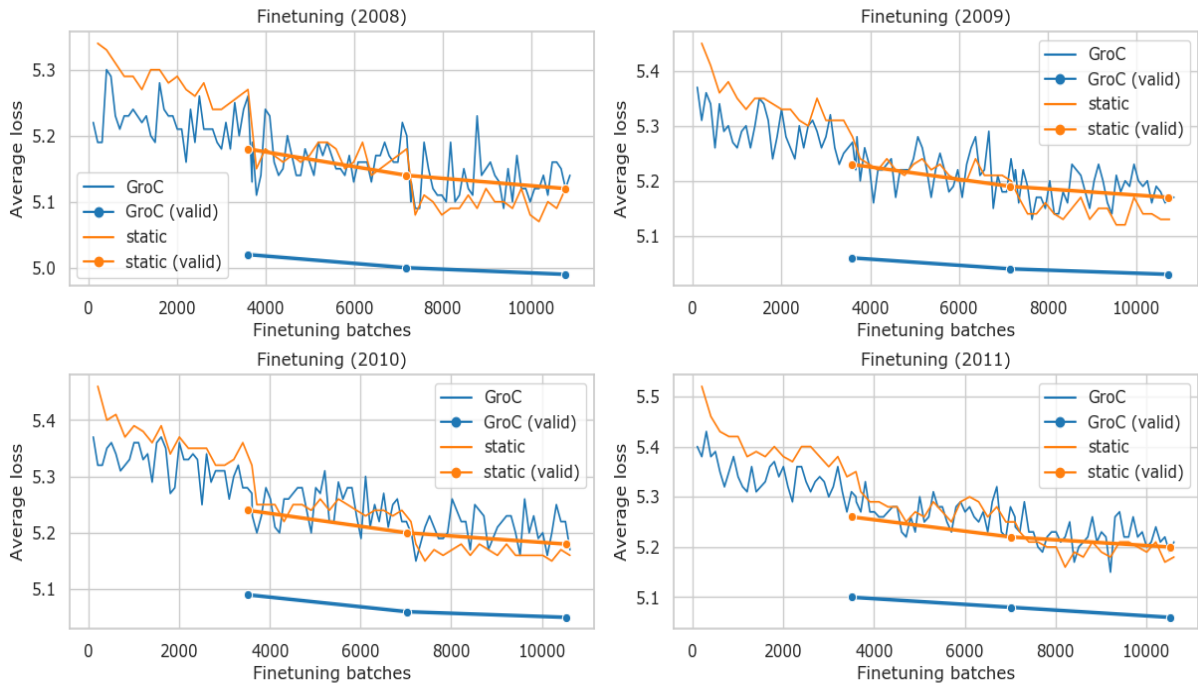


Figure 4: Training and validation loss for GroC and the tied model during finetuning on near domains.

Model	$ \Theta $	test
Mikolov and Zweig (2012) – RNN-LDA	9M [‡]	92.0
Zaremba et al. (2014) – LSTM	20M	82.7
Gal and Ghahramani (2016) – Var. LSTM	20M	78.6
Kim et al. (2016) – CharCNN	19M	78.9
Merity et al. (2017) – Pointer Sentinel-LSTM	21M	70.9
Grave et al. (2017c) – LSTM + cont. cache	-	72.1
Inan et al. (2017) – Tied Variational LSTM	24M	73.2
Zilly et al. (2017) – Variational RHN	23M	65.4
Zoph and Le (2016) – NAS Cell	25M	64.0
Merity et al. (2018) – AWD-LSTM	24M	58.8
Ours – LSTM	10M	86.2
Ours – LSTM + GroC (sur,rel,def)	9.7M	69.5

Table 11: Comparison with state-of-the-art models of comparable size to that of Grave et al. (2017a) and Merity et al. (2017) on the penn dataset.

investigating the potential of competing with even higher capacity models is an interesting direction which we hope will be explored in future studies.

C Cross-Domain Language Modeling

For the experiment in cross-domain language modeling, we used the following computing infrastructure: 2 GeForce RTX 2080 Ti and 2 TITAN RTX GPUs to train and finetune our GroC models, and 2 Tesla P100 GPUs to train and finetune the baselines and to perform hyperparameter search.

C.1 Finetuning Dynamics

Figures 4 and 5 show the loss on the training and validation data for the target domain during finetuning. GroC generalizes better from the training to the validation data than the tied model, consistently having lower validation loss. The training loss for GroC consistently starts out lower than that of the tied model, showing that it has less difficulty adapting to the new data, and ends up higher, indicating greater regularization vs the tied model.

The web dataset is a clear outlier, in which the tied model improves much more dramatically than in any other domain. The difference in validation performance here is reflected in the test perplexity (Table 5) but does not have a clear explanation.

C.2 Data

As described in Section 5, the choice of data and preprocessing used for the cross-domain experiments are based on Grave et al. (2017a). News Crawl and Common Crawl can be downloaded from the WMT 2014 website.¹² WikiText-103 was downloaded from Salesforce website¹³. For the News Crawl datasets, the first 2M tokens of the English data for each year were used as the train

¹²www.statmt.org/wmt14/translation-task.html

¹³blog.einstein.ai/the-wikitext-long-term-dependency-language-modeling-dataset/

Model	2007 →	near domains				far domains	
		2008	2009	2010	2011	Web	Wiki
Tied + finetuning	–	167.44	175.95	177.46	180.63	144.13	232.06
Grounded + finetuning	–	146.84	152.29	155.27	158.21	212.99	188.25

Table 12: Validation perplexity for finetuned models on cross-domain language modeling.




Figure 5: Training and validation loss for GroC and the tied model during finetuning, on far domains.

set, the next 2M tokens as the validation set, and the next 10M tokens as the test set. The same procedure was used for *web* (Common Crawl), for which we used the English portion of the English-German aligned data. While Grave et al. (2017a) describes the Common Crawl data as shuffled at the sentence level, we found that most sentences seemed closely related to adjacent sentences, so after creating train/valid/test splits for this dataset we re-shuffled each file. WikiText-103 comes divided into train/valid/test splits, so we used the first 2M/2M/10M tokens of each split respectively for our dataset. All data was then tokenized using the Europarl tokenizer¹⁴ and lowercased.

Our data preprocessing can be replicated with the script `create-data.sh`, available with the code for GroC.¹⁵

¹⁴statmt.org/europarl/v7/tools.tgz

¹⁵github.com/<anon>/groc

C.3 Finetuning Validation Results

Because no target-domain training is required for most of our cross-domain experiments, validation scores were not computed for most model-domain combinations; however, we report the validation perplexity for the finetuned models in Table 12, to aid in replication.

C.4 Hyperparameter Selection

Cache hyperparameters were selected via grid search, with θ , the flattening hyperparameter described in Grave et al. (2017c), ranging over 5 values from 0 to 1, and λ ranging over 5 values from 0.833 to 0.966 (bounds which were selected based on the optimal hyperparameter ranges in (Grave et al., 2017c)). Perplexity of a model trained on 2007 and evaluated on the 2008 validation set was the metric used to select the optimal hyperparameters: $\lambda = 0.966$ for unigram and neural cache and $\theta = 0.5$ for neural cache. Because the cache is only used during evaluation, this hyperparameter search was quite efficient to carry out using the tied model, requiring no additional training, only 25 evaluation runs on the validation set. This hyperparameter search is illustrated in Figure 6.




Figure 6: Validation accuracy for various hyperparameter settings on the 2008 validation set.

We then used the same hyperparameters for all cache models. This provides a slight advantage to the tied model, as the optimal hyperparameters for GroC might be different from those selected with the tied model. A cache size of 5,000 was used




Figure 7: Validation accuracy for various hyperparameter settings on the `penn` validation set.

during hyperparameter tuning, but at test time we used 10,000 for all experiments based on its use in Grave17. Figure 7 shows a separate hyperparameter search performed over the `penn` validation set to confirm the accuracy of our neural cache reimplementation. Compare to Figure 2a in Grave et al. (2017c); note their λ is 1 minus ours.

For GroC, we also selected a downweighting hyperparameter dw , based on validation performance on the `wiki` dataset only. We searched over 5 values (0.1, 0.3, 0.5, 0.7, and 0.9) using GroC with the neural cache, and selected $dw = 0.1$ as the best value with a validation ppl of 154.01.