

A Cognitive Regularizer for Language Modeling

Jason Wei¹ Clara Meister² Ryan Cotterell^{1,2}
¹Google AI Language ²ETH Zürich University of Cambridge
jasonwei@google.com clara.meister@inf.ethz.ch
ryan.cotterell@inf.ethz.ch

Abstract

The uniform information density (UID) hypothesis, which posits that speakers behaving optimally tend to distribute information uniformly across a linguistic signal, has gained traction in psycholinguistics as an explanation for certain syntactic, morphological, and prosodic choices. In this work, we explore whether the UID hypothesis can be operationalized as an inductive bias for statistical language modeling. Specifically, we augment the canonical MLE objective for training language models with a regularizer that encodes UID. In experiments on ten languages spanning five language families, we find that using UID regularization consistently improves perplexity in language models, having a larger effect when training data is limited. Moreover, via an analysis of generated sequences, we find that UID-regularized language models have other desirable properties, e.g., they generate text that is more lexically diverse. Our results not only suggest that UID is a reasonable inductive bias for language modeling, but also provide an alternative validation of the UID hypothesis using modern-day NLP tools.

1 Introduction

Language has been hypothesized to follow certain information-theoretic constraints. One of the most famous of these constraints is the uniform information density (UID) hypothesis (Fenk and Fenk, 1980; Jaeger, 2010), which states that, subject to the rules of the grammar, speakers aim to distribute information density across a linguistic signal as uniformly as possible. That is, speakers behaving optimally should structure their utterances such that the differences between the peaks and troughs in information are minimized.

In the psycholinguistics literature, the UID hypothesis has been used to explain a variety of linguistic phenomena ranging from how we shorten the phonetic duration of more-predictable linguistic




Figure 1: Graphical illustration of two examples regarding UID. In (a), many speakers will prefer the version with the relativizer *that* (dotted blue line). The UID hypothesis posits that this is because, without the relativizer, the first word of the relative clause, *we*, has high information density; and so including the relativizer distributes the per-word information density more uniformly. In (b), the relativizer *that* is often omitted because, at the onset of the relative clause, the information density of *I* is lower and therefore the distribution of information density is already relatively uniform. Illustration based on Jaeger (2010).

units (Aylett and Turk, 2004) to when we decide to use optional syntactic relativizers (Levy and Jaeger, 2007), among other phenomena (Bell et al., 2003; Frank and Jaeger, 2008). These studies often use language models to estimate the information density of linguistic units, taking observations of low variation of information density in well-formed utterances as evidence for the UID hypothesis.

In this paper, we propose a new experimental paradigm that uses modern-day NLP models to test the UID hypothesis. Whereas prior work has used language modeling as a tool for observing UID,¹ we explore the converse—can UID be used as a tool to train better language models? Specifically, if the UID hypothesis is true, then we should be able to operationalize UID as a regularizer to help train language models. Moreover, observing lower perplexity in language models trained with this regularization would imply that the concept of UID is a good inductive bias for language modeling, thereby providing a new type of evidence for the UID hypothesis *at scale*.

In experiments, we indeed find such evidence: across a variety of languages and dataset sizes, UID regularization consistently improves performance, having a larger effect when training data is limited. Moreover, we observe that—in comparison with their unregularized counterparts—UID-regularized language models are (1) higher entropy while achieving the same (or better) test set perplexity and (2) generate text that is longer and more lexically diverse. Our work is the first to explore the interaction between UID and *training* modern-day neural language models, and our findings—that a cognitively motivated objective can improve language model performance—open up new avenues for testing other psycholinguistic hypotheses in a similar framework.

2 Preliminaries: Language Modeling

The task of language modeling aims to estimate a model of the probability of observing any given string in (a subset of) natural language. Formally, a language model p is an (unconditional) probability distribution over sequences of words $\mathbf{w} = \langle w_1, w_2, \dots \rangle$, where \mathbf{w} consists of tokens from some vocabulary and begins and ends with special tokens BOS and EOS, respectively.

Today’s language models are typically parameterized by neural networks (e.g., transformers (Vaswani et al., 2017)), that follow a local-normalization scheme. Specifically, the model provides a conditional distribution over the vocabulary at each time step; we can then compute the proba-

¹On its own, the term ‘UID’ is formally an *attribute* of a linguistic signal. We also use it throughout this work to refer to the *concept* that UID is a desirable property.

bility of an entire sequence \mathbf{w} as:

$$p_{\theta}(\mathbf{w}) = \prod_{t=1}^{|\mathbf{w}|} p_{\theta}(w_t \mid \mathbf{w}_{<t}) \quad (1)$$

where θ are the parameters of the model and we use $\mathbf{w}_{<t}$ to represent the first $t - 1$ tokens of \mathbf{w} . Parameters are estimated by optimizing over some objective $L(\theta)$. The standard objective for language modeling is the negative log-likelihood of a dataset \mathcal{W} under the model:

$$L(\theta) = - \sum_{\mathbf{w} \in \mathcal{W}} \log p_{\theta}(\mathbf{w}) \quad (2)$$

Subsequently, we drop explicit dependence on θ when it is obvious from context.

To assess the goodness of fit of a model p , we typically evaluate its perplexity on some held-out dataset $\mathcal{W}_{\text{test}}$, where perplexity (PPL) is defined as

$$\text{PPL}(p) = \exp \left(- \sum_{\mathbf{w} \in \mathcal{W}_{\text{test}}} \frac{1}{|\mathbf{w}|} \log p(\mathbf{w}) \right) \quad (3)$$

Note that under this definition of perplexity, our evaluation metric is slightly different than the training objective; the former computes an average over each sequence while the later treats all tokens equally, regardless of the length of the sequence in which they are present.

3 Uniform Information Density

Communication via natural language is a complicated and nuanced process that takes place under a host of cognitive and environmental constraints. As a result, speakers have to make (perhaps subconscious) choices to best navigate this communicative dance. A rational speaker would use these choices to optimize the communicative properties of their utterances. One such locus of optimization is outlined by the **Uniform Information Density (UID)** hypothesis.

3.1 The UID Hypothesis

At its core, the UID hypothesis aims to explain certain phenomena in human language processing using an information-theoretic approach: we can view language as a transfer of information, which is transmitted with a certain density through a communication channel. The UID hypothesis posits that speakers that behave optimally will structure

their utterances to avoid peaks and troughs in this information density (Aylett and Turk, 2004; Levy and Jaeger, 2007; Jaeger, 2010). More formally stated: “*Within the bounds defined by grammar, speakers prefer utterances that distribute information uniformly across the signal (information density). Where speakers have a choice between several variants to encode their message, they prefer the variant with more-uniform information density (ceteris paribus)*” (Jaeger, 2010).

3.2 Example: UID in syntactic reduction

To better understand the UID hypothesis, consider the concrete example of syntactic reduction (*that*-mentioning) from Jaeger (2010), which we show graphically in Figure 1 and also describe below.

Ex. A. My boss confirmed [that] we are crazy.

Ex. B. My boss thinks [that] I am crazy.

In both these sentences, the use of the relativizer *that* is syntactically optional—at the onset of a relative clause (RC), speakers can, but do not have to, include the relativizer. Many speakers, however, would argue that the sentence flows better with the relativizer included in Example A and the relativizer omitted in Example B.

The UID hypothesis provides a potential explanation for this phenomenon. When a RC is used without a relativizer, the first word of the RC conveys two pieces of information: both the onset of the RC, as well as part of the RC’s internal contents. In Example A, many speakers would find that the information density of the first word in the RC, *we*, is high, and so adding in the relative clause distributes the information over two words, making it easier to parse. In Example B, the information density of the first word in the RC, *I*, is lower relatively, and so we do not need to (or it is not as beneficial to) include the relativizer.

3.3 Measuring UID

Now that we better understand what the UID hypothesis attempts to explain, how might we operationalize UID and find quantitative evidence of the pressure for it in language? First, to quantify the amount of information conveyed by a word, we turn to the most basic information-theoretic definition: the information conveyed by a word w in context is its Shannon information content (Shannon, 1948), also called *surprisal*. Ideally, this surprisal would be measured using the “true” distribution over human language. Because we do not have access to

such a distribution, we often estimate it using a statistical language model. That is, given a statistical language model p , which estimates the probability of a word given its context, the surprisal $u(w_t)$ of word w_t is defined as the following:

$$u(w_t) = -\log p(w_t \mid \mathbf{w}_{<t}) \quad (4)$$

This setup provides a natural approach to exploring how UID might manifest—if the UID hypothesis is true, then we should observe that variation in surprisal, as estimated by a language model, is minimized in natural language.

Using this approach, prior work has accumulated evidence for UID across various levels of linguistic representation (Pluymaekers et al., 2005; Bell et al., 2009, *inter alia*). As some of the earliest examples, Aylett and Turk (2004) showed that linguistic units that had high surprisal according to a tri-gram language model were uttered with longer syllable durations, and Levy and Jaeger (2007) found that for RCs in which the first word had higher surprisal, relativizers were more likely to be used in the RC during actual speech. Further examples are given in our related work section (§7).

4 UID-Regularized Language Modeling

While prior work has shown evidence that UID can help explain many of the choices we make when generating language, to the best of our knowledge, operationalizations of UID have not been explicitly employed as part of the training objective in modern-day NLP models. This raises the simple question that is central to our paper:

Can UID serve as an inductive bias for training statistical language models?

In an effort to answer this question, we present a scheme for incorporating operationalizations of UID into the language model training objective. Formally, we augment the canonical maximum likelihood estimation objective² in eq. (2) with UID

²Note that the maximum likelihood estimation objective minimizes (over $\mathbf{w} \in \mathcal{W}$) $-\log p(w_t \mid \mathbf{w}_{<t})$, i.e., surprisal. Although such an objective may indirectly minimize peaks and dips in surprisal across a sequence simply by pushing them towards 0, it does not explicitly include any sequence level penalty for even surprisal distribution.

operationalizations as regularizers \mathcal{R} . Under this new objective, we minimize

$$L_{\mathcal{R}}(\theta) = L(\theta) + \beta \cdot \mathcal{R}(\theta) \quad (5)$$

where $\beta > 0$ is the strength coefficient of the regularizer. We consider two natural operationalizations of UID—inspired by Collins (2014)—as regularizers for training language models:

Variance Regularizer. UID concerns the distribution of information in language production, and so a natural measure of this behavior is the variance of surprisals. Thus, we first consider a regularizer that penalizes high variance among the surprisals of words in a given sequence:

$$\mathcal{R}(\theta) = \frac{1}{|\mathbf{w}|} \sum_{t=1}^{|\mathbf{w}|} (u(w_t) - \mu)^2 \quad (6)$$

where $\mu = \frac{1}{|\mathbf{w}|} \sum_{t=1}^{|\mathbf{w}|} u(w_t)$. Note that here, and in our subsequent regularizers, we estimate $u(\cdot)$ via eq. (4) using our model p_{θ} .

Local Consistency. Next, we consider a local consistency regularizer that encourages the surprisals of adjacent words to have similar magnitude:

$$\mathcal{R}(\theta) = \frac{1}{|\mathbf{w}|-1} \sum_{t=1}^{|\mathbf{w}|-1} \left(u(w_t) - u(w_{t+1}) \right)^2 \quad (7)$$

This regularizer is also a reasonable operationalization of UID—if every surprisal is similar to its neighbor, then the density of information in the sequence will be close to uniform.

Though we focus on these two regularizers, other operationalizations of UID certainly exist. For example, a similar variant of the above regularizers is the max regularizer (Meister et al., 2020a), which penalizes the highest surprisal in a sentence.³ Furthermore, UID may also be defined in terms of parse steps (Hale, 2001) or structural integrations (Gibson, 2000), as well as in spoken language in the form of filler words like *uh* and *um* or word repetition during challenging lexical retrieval. We consider these operationalizations (as well as the broader discussion of how to operationalize UID) as future work.

³We also tried this operationalization in preliminary experiments, but results were not as strong as the variance or local consistency regularizers.

5 Experimental Setup

To empirically evaluate UID regularization, we train various language models with the UID-regularized objective (eq. (5)) using the following experimental setup.

Datasets. We employ datasets from multiple languages and of varying sizes. We use the EuroParl corpus (Koehn, 2005)—a multi-lingual dataset of discussions from the European Parliament that has been commonly used for language modeling (Cotterell et al., 2018; Mielke et al., 2019)—since it is roughly semantically controlled in that all utterances are presumably about the same topics. We use EuroParl v7 download from the ACL 2014 SMT Workshop⁴ and perform a 80–10–10 train-dev-test split on all five languages—Czech, English, French, German, and Spanish—which yields 46.7, 42.2, 47.2, 51.3, and 12.4 million training tokens for each language respectively.

Moreover, we experiment on languages from several language families; the five languages in EuroParl that we consider are all Indo-European, and so we look to Wiki-40B (Guo et al., 2020), which contains Wikipedia dumps of a wide range of languages. We choose a set of diverse languages with training set sizes relatively similar to that of EuroParl: Finnish (a Uralic language; 59.3M training tokens), Indonesian (an Austronesian language; 45.7M training tokens), and Turkish (a Turkic language; 38.1M training tokens). To explore performance on lower-resource languages, we additionally experiment with Swahili⁵ (a Niger-Congo language; 6.3M training tokens) and Tagalog (an Austronesian language; 4.2M training tokens). For all languages, we performed tokenization using the MosesTokenizer.⁶ Train, dev, and test set splits are shown in Table 5 in the Appendix.

Model Framework and Architecture. For our experiments, we use the fairseq library (Ott et al., 2019), a standard sequence modeling toolkit in PyTorch. As our model, we use fairseq’s default transformer (with six decoder layers and eight

⁴<http://statmt.org/wmt14/translation-task.html>

⁵Since there are no Niger-Congo languages in Wiki-40B, we perform a 80-10-10 split on Swahili Wikidumps (see <https://github.com/google-research/bert/blob/master/multilingual.md>).

⁶<https://pypi.org/project/mosestokenizer/>

attention heads), which achieves competitive⁷ language modeling performance (although the purpose of our paper is not to achieve or compare with the state of the art). For all experiments, we followed the data-preprocessing scripts and recommended hyperparameters provided in `fairseq`'s language modeling module; more detailed information can be found on the Github page.⁸

UID Regularizers. For UID regularization, we experiment with the variance (eq. (6)) and local consistency regularizers (eq. (7)). We found in preliminary experiments that effective regularization strengths were often near $\beta = 0.01$, and so we performed a grid search over values within an order of magnitude around $\beta = 0.01$: $\beta \in \{0.006, 0.008, 0.01, 0.02, 0.03, 0.04, 0.05\}$. We choose the model with the lowest dev loss to evaluate on the test set.

6 Results

In this section, we report results for models trained under the UID-regularized objective. We find that UID regularization consistently improves perplexity for models trained on various languages (§6.1) and dataset sizes (§6.2). Additionally, we examine properties of text generated by UID-regularized models (§6.3) and analyze the relationship between our operationalization of UID and perplexity (§6.4).

6.1 Languages

Table 1 shows the results of UID-regularized language models trained on various languages from EuroParl and Wiki-40B, and includes statistical significance of changes in perplexity, as compared with baselines, computed using permutation tests⁹ (Efron and Tibshirani, 1994). For all languages, UID regularization significantly improves perplexity for at least one of the two regularizers. Further-

⁷On Wikitext-103, the largest dataset we train on (103 million tokens), we achieve a competitive perplexity of 29.89 (c.f. Merity et al. (2018)). For smaller datasets, we tried a smaller transformer architecture of four decoder layers and four attention heads, but it did not perform better than the six decoder layer and eight attention heads version, suggesting that this architecture was not too large for the datasets we use in this paper (even the Tagalog dataset we use is larger than the commonly used Penn Treebank and WikiText-2).

⁸https://github.com/pytorch/fairseq/tree/master/examples/language_model

⁹<http://www2.stat.duke.edu/~ar182/rr/examples-gallery/PermutationTest.html>

LANGUAGE (# train tokens)	Perplexity
CZECH (12.4M)	
Baseline (no UID)	47.47
+ UID: variance	47.24 (↓0.5%) [†]
+ UID: local consistency	47.08 (↓0.8%) [†]
ENGLISH (46.7M)	
Baseline (no UID)	21.34
+ UID: variance	21.08 (↓1.2%) [†]
+ UID: local consistency	21.19 (↓0.7%) [†]
FINNISH (59.3M)	
Baseline (no UID)	51.58
+ UID: variance	51.30 (↓0.5%) [†]
+ UID: local consistency	51.49 (↓0.2%)
FRENCH (51.3M)	
Baseline (no UID)	17.08
+ UID: variance	17.02 (↓0.4%) [†]
+ UID: local consistency	17.03 (↓0.3%) [†]
GERMAN (42.3M)	
Baseline (no UID)	26.62
+ UID: variance	26.50 (↓0.4%) [†]
+ UID: local consistency	26.45 (↓0.6%) [†]
INDONESIAN (45.7M)	
Baseline (no UID)	53.96
+ UID: variance	53.66 (↓0.6%) [†]
+ UID: local consistency	53.70 (↓0.5%)
SPANISH (47.2M)	
Baseline (no UID)	22.54
+ UID: variance	22.37 (↓0.8%) [†]
+ UID: local consistency	22.44 (↓0.4%) [†]
SWAHILI (6.3M)	
Baseline (no UID)	40.45
+ UID: variance	39.79 (↓1.6%) [†]
+ UID: local consistency	39.44 (↓2.5%) [†]
TAGALOG (4.2M)	
Baseline (no UID)	80.48
+ UID: variance	78.40 (↓2.5%) [†]
+ UID: local consistency	78.12 (↓2.9%) [†]
TURKISH (38.1M)	
Baseline (no UID)	66.13
+ UID: variance	65.70 (↓0.7%) [†]
+ UID: local consistency	66.06 (↓0.1%)

Table 1: UID regularizers improve perplexity for multiple languages. [†] indicates statistical significance compared with the baseline ($p < 0.05$).

more, UID regularization (under the best performing β) never leads to worse perplexity. These results suggest that incorporating UID operationalizations into a model's training objective leads to a better model of language, substantiating uniform information density as a valid inductive bias. Moreover, the improvement for many languages corroborates the expectation that UID should, due to its information theoretic nature, hold across languages (Jaeger and Tily, 2011).

	WMT'06	EuroParl	WT-103
# training tokens	16.0M	47.0M	103.2M
Baseline (no UID)	49.70	21.34	29.89
+ UID: variance	48.25[†]	21.08[†]	29.58
+ UID: local consistency	48.79	21.19	29.73

Table 2: UID regularizers improve perplexity on language models trained on English datasets of varying size. Improvements tend to be larger on smaller datasets. [†] indicates statistical significance compared with the baseline ($p < 0.05$).

6.2 Dataset Size

Notably, we observe the largest improvements (1.6–2.9%) in perplexity in Table 1 for the lowest resource languages, Tagalog and Swahili (with 4.2 and 6.3 million training tokens respectively). Conversely, improvement was most marginal (0.2–0.5%) on the highest-resource languages, French and Finnish (51.3 and 59.3 million training tokens respectively). To remove language as a confounding factor from this observation, we perform a controlled analysis of the effects of UID regularization as a function of dataset size.

We focus on English; in addition to the result on English EuroParl 2014 from Table 1, which contains 47.0 million training tokens, we experiment with the smaller monolingual English dataset from the 2006 NAACL Workshop on Statistical Machine Translation (WMT'06),¹⁰ which has 17.0M tokens in its training set, as well as the larger Wikitext-103 benchmark (Merity et al., 2017), which contains 103 million tokens in its training set.

Table 2 shows the perplexities for models with and without UID regularization for these three datasets. As suggested by earlier results, improvements were strongest for the WMT'06 dataset, with an improvement of 1.4 perplexity points for the variance regularizer and 0.9 PPL points for local consistency. For the larger EuroParl and WT-103 datasets, on the other hand, improvement was more modest, ranging from 0.1 to 0.3 perplexity points.

As further confirmation that UID regularization has a greater impact on smaller datasets, we perform an ablation study that roughly controls for language content by training models on the subsets of the same dataset. For this ablation, we take subsets of 2, 4, 8, 12, 16, 24, and 32 million sentences from the 47 million sentences in English EuroParl,

¹⁰We downloaded the given train-dev-test splits from <https://www.statmt.org/wmt06/>.




Figure 2: Improvement in perplexity for UID regularized models trained on subsets of varying size sampled from the EuroParl English dataset (full dataset size 47.0 million tokens). UID regularization helped more when training data was more limited.

and observe how much the UID regularizers improve perplexity for each training dataset size. As shown in Figure 2, the results tell the same story as Table 2—UID regularization improves perplexity more for smaller datasets.

These results are consistent with the expectation that models trained on smaller datasets are more likely to overfit and could therefore benefit more from regularization (Melis et al., 2018). As it is possible that the models trained on smaller datasets could benefit from any kind of regularization, we experiment with label smoothing (Szegedy et al., 2016), another regularization technique that similarly augments the training objective with a penalty. Table 4 shows these results for models trained on WMT'06 and EuroParl with label smoothing—our experiments indicate that, across the board, label smoothing leads to worse perplexity compared with baseline models.¹¹ We take this result as further evidence that the improvement from UID regularization stems from the UID hypothesis as a valid inductive bias, rather than simply a need for any kind of regularization when training on smaller datasets.

¹¹This negative result for applying label smoothing to language modeling is consistent with prior empirical findings (Müller et al., 2019; Gao et al., 2020; Meister et al., 2020b).

	Sequence length	Model entropy	% unique n -grams		
			$n = 2$	$n = 3$	$n = 4$
Baseline (no UID)	22.9	69.6	37.7	73.5	90.9
+ UID: variance	24.0	79.4	40.7	77.8	93.3
+ UID: local consistency	23.3	73.9	39.1	75.7	92.1

Table 3: Text generated by UID-regularized language models is longer (higher average sequence length), higher entropy (computed via monte-carlo estimation), and more lexically diverse (a higher ratio of unique n -grams).

	WMT'06	EuroParl
# training tokens	16.0M	47.0M
Baseline	35.75	23.22
+ label smoothing, $\alpha = 0.01$	36.15	26.26
+ label smoothing, $\alpha = 0.05$	55.56	40.79
+ label smoothing, $\alpha = 0.1$	90.57	68.26

Table 4: Label smoothing, another form of regularization that similarly augments the cross-entropy objective with a penalty, does not improve perplexity. (Results shown on dev set).

6.3 Evaluating Generated Text

Unconditional models of language have been observed to produce generic text that can be short, bland, or repetitive (Fan et al., 2018; Kulikov et al., 2019; Holtzman et al., 2020), and so in this subsection we investigate how UID regularization might affect these characteristics in generated text. For these experiments, we consider the baseline model, the variance-regularized model, and the local consistency-regularized model trained on English EuroParl. To obtain text samples, we generate samples by sequentially sampling tokens according to the model’s predicted distribution until the end-of-sequence (EOS) token is sampled, i.e., ancestral sampling. Note that for language model p , this sampling scheme is equivalent to directly sampling $\mathbf{y} \sim p$. We obtain 10,000 samples for each model and report statistics in Table 3.

We analyze each set of generated sentences for several metrics. First, we compute the average length of generated sentences. Next, we evaluate the lexical diversity of generated texts by computing the percent of unique n -grams for $n \in \{2, 3, 4\}$. Finally, sampling from a model also gives us a means for estimating the language model’s entropy:

$$H(p) = - \sum_{\mathbf{y} \in \text{supp}(p)} p(\mathbf{y}) \log p(\mathbf{y}) \quad (8)$$

$$= -\mathbb{E}_{\mathbf{y} \sim p} (\log p(\mathbf{y})) \quad (9)$$

In the case of language models, $\text{supp}(p)$ is the set of all strings that can be generated from the model’s

vocabulary \mathcal{V} . As this is exponentially large in $|\mathcal{V}|$, directly computing $H(p)$ is intractable. We can use its equivalence to eq. (9), however, to estimate $H(p)$ with a simple Monte-Carlo estimator:

$$\hat{H}(p) = -\frac{1}{K} \sum_{k=1}^K \log p(\mathbf{y}^{(k)}) \quad (10)$$

where we sample $\mathbf{y}^{(k)} \sim p$ for $k = 1, \dots, K$.

Table 3 shows results from UID-regularized models compared with the baseline. The models trained with the variance and local consistency regularizers exhibit a preference for longer sequence length and higher lexical diversity. Additionally, the entropy estimates of these models are notably higher, which, following the principle of maximum entropy (Jaynes, 1957),¹² can be seen as an additional advantage of UID-regularized models over their unregularized counterparts.

6.4 UID Behavior

To take a closer look at how UID regularization affects language models, we examine the relationship between minimizing perplexity and UID behavior, where we quantify UID behavior as the variance of models’ surprisals. We consider models trained on the English EuroParl dataset with the variance regularizer at strengths $\beta \in \{0.01, 0.03, 0.05, 0.07, 0.09\}$ and our baseline (which is equivalent to $\beta = 0$). For further comparison, we also train a model with $\beta = -0.01$ to observe the effects of *penalizing* UID behavior. We report results on the EuroParl test set in Figure 3.

We observe that the model trained with a UID penalty (negative β) indeed exhibits worse perplexity and UID behavior (variance of surprisals) on the test set. And as we might expect, models trained with higher β exhibit UID behavior more strongly, as our quantification is part of their training objective. Overall, from $\beta = 0.01$ to $\beta = 0.05$, both

¹²The principle of maximum entropy states that the probability distribution that best represents the current knowledge state is the one with the largest entropy.

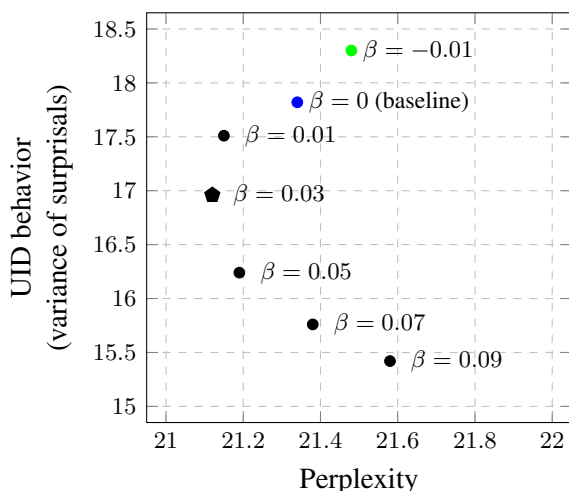


Figure 3: A trade-off between perplexity (x -axis) and variance of surprisals (a measure of UID behavior; y -axis). The black pentagon indicates the β that yielded the best perplexity ($\beta = 0.03$).

perplexity and UID behavior are positively correlated with β , but when we optimize too much for UID ($\beta \geq 0.07$), there is a trade-off in which model perplexity begins to increase.

We also observe an intriguing phenomenon in Figure 3. Models that achieve similar perplexity can have substantially different UID behavior values on the test set. Specifically, the $\beta = 0$ and $\beta = 0.07$ models, which have almost the same perplexity, have variance of surprisals of 17.8 and 15.8—a difference of more than ten percent! If such models with similar perplexity can have varying definitions of what constitutes good UID behavior, then prior work, which has drawn conclusions on UID based on surprisals computed by a single model (Aylett and Turk, 2004; Levy and Jaeger, 2007; Jain et al., 2018), may need revisiting. As this direction is outside the scope of the present paper, we leave it as future work.

7 Discussion and Related Work

We discussed how operationalizing UID for language modeling leads to better models in a wide variety of settings. These results both provide a new form of evidence for the UID hypothesis and build on prior work exploring UID in modern-day NLP models.

Evidence for the UID hypothesis. Our work extends the body of psycholinguistic research on uniform information density, which has largely corroborated the UID hypothesis by providing evidence that variation in surprisal, as estimated by a lan-

guage model, is minimized in natural language. In addition to early studies that used this approach to find evidence for UID in syntactic reduction (Levy and Jaeger, 2007), morphosyntactic contractions (Frank and Jaeger, 2008), and prosodic structure (Aylett and Turk, 2004), the same line of reasoning has been used by more recent work exploring a variety of other linguistic properties. These studies have found that word duration can be predicted by syntactic surprisal (Demberg et al., 2012; Moore-Cantwell, 2013), construction probability (Kuperman and Bresnan, 2012), informativity (Seyfarth, 2014), and contextual predictability (Jurafsky et al., 2001; Bell et al., 2003; Gahl and Garnsey, 2004). They have also observed that word length is reflected by conceptual complexity (Lewis and Frank, 2016); word order choice can be predicted by processing cost (Bloem, 2016; Sikos et al., 2017); phonological patterns can be shaped by word predictability (Hall et al., 2018); and UID computed at the sequence level predicts human preferences for syntactic alternatives of the same sentence.

Whereas the above prior work has used language modeling as a tool for measuring UID, our paper has explored the exact converse—we have asked whether UID, operationalized as a regularizer, can be used as a tool for training better language models. We argue that if the UID hypothesis holds as a general principle, then we should be able to exploit it as a training criterion that improves language modeling. And accordingly, our results show that—across a variety of languages and dataset sizes—regularization for UID did indeed improve perplexity, which we view as an alternative kind of evidence for the UID hypothesis at scale.

Notably, Figure 3 at first could appear to contradict the UID hypothesis, since models with better UID behavior did not always achieve better perplexity. We do not consider this as evidence against the UID hypothesis, however. Rather, we posit that when β is too large, we may be optimizing for UID to the point of tending towards unnatural language—a perfectly uniform dispersion of information across an utterance may come at the cost of strange lexical choices. In this light, such a trade-off should be somewhat expected.

UID in modern NLP. In addition to the traditional line of psycholinguistic work, there have also been more-recent studies on UID in the context of modern NLP, although this work is relatively sparse. Rubino et al. (2016) leverage infor-

mation density encoded as surprisal at the word, part of speech, and syntax levels to help build a state-of-the-art model for mixed-domain translationese detection. Jain et al. (2018) incorporate UID measures across sentences into models designed to detect natural versus manipulated text. Perhaps the work that is most related to ours, Meister et al. (2020a), leverages UID to explain why beam search is an effective decoding algorithm and uses operationalizations of UID during beam search to alleviate problems with decoding poorly calibrated machine translation models. Whereas Meister et al. (2020a) focuses on decoding, our work shows the first evidence that UID can be operationalized to aid training.

8 Conclusions

In closing, we have proposed encoding uniform information density as a regularizer for training language models—a novel manner of incorporating an established psycholinguistic theory into modern statistical language modeling. In experiments on a range of languages and dataset sizes, UID regularization consistently improves perplexity over baselines. Our results suggest that UID is a valid inductive bias for improving the canonical maximum likelihood objective in language modeling, providing a new, alternative type of evidence that supports the UID hypothesis at scale. Our work opens the door to future research directions such as using similar techniques to validate other psycholinguistic phenomena, applying UID regularization in conditional language generation tasks, and exploring how UID regularized models perform in downstream NLP applications.

Ethical Concerns

Language models have various ethical, environmental, and financial concerns. We cannot do justice to them here, but do see Bender et al. (2021) for a pointer. We do not foresee any additional ethical concerns with the contributions made in our work beyond those discussed in Bender et al. (2021).

Acknowledgements

We thank Roger Levy for feedback in the middle stages of our work and Tiago Pimentel, David Reitter, Tal Linzen, and Slav Petrov for feedback on the manuscript.

References

- Matthew Aylett and Alice Turk. 2004. [The smooth signal redundancy hypothesis: A functional explanation for relationships between redundancy, prosodic prominence, and duration in spontaneous speech](#). *Language and Speech*, 47(1):31–56. PMID: 15298329.
- Alan Bell, Jason M. Brenier, Michelle Gregory, Cynthia Girand, and Dan Jurafsky. 2009. [Predictability effects on durations of content and function words in conversational English](#). *Journal of Memory and Language*, 60(1):92–111.
- Alan Bell, Daniel Jurafsky, Eric Fosler-Lussier, Cynthia Girand, Michelle Gregory, and Daniel Gildea. 2003. [Effects of disfluencies, predictability, and utterance position on word form variation in English conversation](#). *The Journal of the Acoustical Society of America*, 113(2):1001–1024.
- Emily M. Bender, Timnit Gebru, Angelina McMillan-Major, and Shmargaret Shmitchell. 2021. [On the dangers of stochastic parrots: Can language models be too big?](#) In *Proceedings of the 2021 ACM Conference on Fairness, Accountability, and Transparency, FAccT '21*, page 610–623, New York, NY, USA. Association for Computing Machinery.
- Jelke Bloem. 2016. [Testing the processing hypothesis of word order variation using a probabilistic language model](#). In *Proceedings of the Workshop on Computational Linguistics for Linguistic Complexity (CLALC)*, pages 174–185, Osaka, Japan. The COLING 2016 Organizing Committee.
- Michael Xavier Collins. 2014. [Information density and dependency length as complementary cognitive models](#). *Journal of Psycholinguistic Research*, 43(5):651–681.
- Ryan Cotterell, Sabrina J. Mielke, Jason Eisner, and Brian Roark. 2018. [Are all languages equally hard to language-model?](#) In *Proceedings of the 2018 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies, Volume 2 (Short Papers)*, pages 536–541, New Orleans, Louisiana. Association for Computational Linguistics.
- Vera Demberg, Asad Sayeed, Philip Gorinski, and Nikolaos Engonopoulos. 2012. [Syntactic surprisal affects spoken word duration in conversational contexts](#). In *Proceedings of the 2012 Joint Conference on Empirical Methods in Natural Language Processing and Computational Natural Language Learning*, pages 356–367, Jeju Island, Korea. Association for Computational Linguistics.
- Bradley Efron and Robert J. Tibshirani. 1994. *An Introduction to the Bootstrap*. CRC Press.
- Angela Fan, Mike Lewis, and Yann Dauphin. 2018. [Hierarchical neural story generation](#). In *Proceedings of the 56th Annual Meeting of the Association for*

- Computational Linguistics (Volume 1: Long Papers)*, pages 889–898, Melbourne, Australia. Association for Computational Linguistics.
- August Fenk and Gertraud Fenk. 1980. [Konstanz im kurzzeitgedächtnis-konstanz im sprachlichen informationsfluß](#). *Zeitschrift für Experimentelle und Angewandte Psychologie*, 27:400–414.
- Austin F. Frank and T. Florian Jaeger. 2008. [Speaking rationally: Uniform information density as an optimal strategy for language production](#). In *Proceedings of the Annual Meeting of the Cognitive Science Society*, volume 30.
- Susanne Gahl and Susan M. Garnsey. 2004. [Knowledge of grammar, knowledge of usage: Syntactic probabilities affect pronunciation variation](#). *Language*, pages 748–775.
- Yingbo Gao, Weiyue Wang, Christian Herold, Zijian Yang, and Hermann Ney. 2020. [Towards a better understanding of label smoothing in neural machine translation](#). In *Proceedings of the 1st Conference of the Asia-Pacific Chapter of the Association for Computational Linguistics and the 10th International Joint Conference on Natural Language Processing*, pages 212–223, Suzhou, China. Association for Computational Linguistics.
- Edward Gibson. 2000. [The dependency locality theory: A distance-based theory of linguistic complexity](#). *Image, language, brain: Papers from the first mind articulation project symposium*, 2000:95–126.
- Mandy Guo, Zihang Dai, Denny Vrandečić, and Rami Al-Rfou. 2020. [Wiki-40B: Multilingual language model dataset](#). In *Proceedings of the 12th Language Resources and Evaluation Conference*, pages 2440–2452, Marseille, France. European Language Resources Association.
- John Hale. 2001. [A probabilistic Earley parser as a psycholinguistic model](#). In *Second Meeting of the North American Chapter of the Association for Computational Linguistics*.
- Kathleen Currie Hall, Elizabeth Hume, T. Florian Jaeger, and Andrew Wedel. 2018. [The role of predictability in shaping phonological patterns](#). *Linguistics Vanguard*, 4(s2).
- Ari Holtzman, Jan Buys, Maxwell Forbes, and Yejin Choi. 2020. [The curious case of neural text degeneration](#). In *Proceedings of the International Conference on Learning Representations*.
- T. Florian Jaeger. 2010. [Redundancy and reduction: Speakers manage syntactic information density](#). *Cognitive Psychology*, 61(1).
- T. Florian Jaeger and Harry Tily. 2011. [On language ‘utility’: Processing complexity and communicative efficiency](#). *Wiley Interdisciplinary Reviews: Cognitive Science*, 2.
- Ayush Jain, Vishal Singh, Sidharth Ranjan, Rajakrishnan Rajkumar, and Sumeet Agarwal. 2018. [Uniform Information Density effects on syntactic choice in Hindi](#). In *Proceedings of the Workshop on Linguistic Complexity and Natural Language Processing*, pages 38–48, Santa Fe, New-Mexico. Association for Computational Linguistics.
- Edwin T. Jaynes. 1957. [Information Theory and Statistical Mechanics](#). *Physical Review*, 106(4):620.
- Daniel Jurafsky, Alan Bell, Michelle Gregory, and William D. Raymond. 2001. [Probabilistic relations between words: Evidence from reduction in lexical production](#). *Typological Studies in Language*, 45:229–254.
- Philipp Koehn. 2005. [Europarl: A parallel corpus for statistical machine translation](#). In *MT Summit*, pages 79–86.
- Iliia Kulikov, Alexander Miller, Kyunghyun Cho, and Jason Weston. 2019. [Importance of search and evaluation strategies in neural dialogue modeling](#). In *Proceedings of the 12th International Conference on Natural Language Generation*, pages 76–87, Tokyo, Japan. Association for Computational Linguistics.
- Victor Kuperman and Joan Bresnan. 2012. [The effects of construction probability on word durations during spontaneous incremental sentence production](#). *Journal of Memory and Language*, 66(4):588–611.
- Roger P. Levy and T. F. Jaeger. 2007. [Speakers optimize information density through syntactic reduction](#). In *Advances in Neural Information Processing Systems*.
- Molly L. Lewis and Michael C. Frank. 2016. [The length of words reflects their conceptual complexity](#). *Cognition*, 153:182–195.
- Clara Meister, Ryan Cotterell, and Tim Vieira. 2020a. [If beam search is the answer, what was the question?](#) In *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing (EMNLP)*, pages 2173–2185, Online. Association for Computational Linguistics.
- Clara Meister, Elizabeth Salesky, and Ryan Cotterell. 2020b. [Generalized entropy regularization or: There’s nothing special about label smoothing](#). In *Proceedings of the 58th Annual Meeting of the Association for Computational Linguistics*, Online. Association for Computational Linguistics.
- Gábor Melis, Chris Dyer, and Phil Blunsom. 2018. [On the state of the art of evaluation in neural language models](#). In *Proceedings of the International Conference on Learning Representations*.
- Stephen Merity, Nitish Shirish Keskar, and Richard Socher. 2018. [An analysis of neural language modeling at multiple scales](#). *CoRR*, abs/1803.08240.

- Stephen Merity, Caiming Xiong, James Bradbury, and Richard Socher. 2017. [Pointer sentinel mixture models](#). In *Proceedings of the International Conference on Learning Representations*.
- Sabrina J. Mielke, Ryan Cotterell, Kyle Gorman, Brian Roark, and Jason Eisner. 2019. [What kind of language is hard to language-model?](#) In *Proceedings of the 57th Annual Meeting of the Association for Computational Linguistics*, pages 4975–4989, Florence, Italy. Association for Computational Linguistics.
- Claire Moore-Cantwell. 2013. [Syntactic predictability influences duration](#). In *Proceedings of Meetings on Acoustics*. Acoustical Society of America.
- Rafael Müller, Simon Kornblith, and Geoffrey E. Hinton. 2019. [When does label smoothing help?](#) In *Advances in Neural Information Processing Systems*.
- Myle Ott, Sergey Edunov, Alexei Baevski, Angela Fan, Sam Gross, Nathan Ng, David Grangier, and Michael Auli. 2019. [fairseq: A fast, extensible toolkit for sequence modeling](#). In *Proceedings of the 2019 Conference of the North American Chapter of the Association for Computational Linguistics (Demonstrations)*, pages 48–53, Minneapolis, Minnesota. Association for Computational Linguistics.
- Mark Pluymaekers, Mirjam Ernestus, and R. Harald Baayen. 2005. [Lexical frequency and acoustic reduction in spoken dutch](#). *The Journal of the Acoustical Society of America*, 118(4):2561–2569.
- Raphael Rubino, Ekaterina Lapshinova-Koltunski, and Josef van Genabith. 2016. [Information density and quality estimation features as translationese indicators for human translation classification](#). In *Proceedings of the 2016 Conference of the North American Chapter of the Association for Computational Linguistics: Human Language Technologies*, pages 960–970, San Diego, California. Association for Computational Linguistics.
- Scott Seyfarth. 2014. [Word informativity influences acoustic duration: Effects of contextual predictability on lexical representation](#). *Cognition*, 133(1):140–155.
- Claude E. Shannon. 1948. [A mathematical theory of communication](#). *The Bell System Technical Journal*, 27(3):379–423.
- Les Sikos, Clayton Greenberg, Heiner Drenhaus, and Matthew W. Crocker. 2017. [Information density of encodings: The role of syntactic variation in comprehension](#). In *Proceedings of the 39th Annual Meeting of the Cognitive Science Society*.
- Christian Szegedy, Vincent Vanhoucke, Sergey Ioffe, Jon Shlens, and Zbigniew Wojna. 2016. [Rethinking the inception architecture for computer vision](#). In *Proceedings of the IEEE Conference on Computer Vision and Pattern Recognition*, pages 2818–2826.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. [Attention is all you need](#). In *Advances in Neural Information Processing Systems*.

A Appendix

Datasets. Table 5 shows the train, dev, and test set splits for the language modeling datasets we use.

Language	Family	Source	Split	Vocab size	Train		Dev		Test	
					Sentences	Tokens	Sentences	Tokens	Sentences	Tokens
English	Indo-European	EuroParl	80–10–10	64k	1.6M	46.7M	201k	5.8M	201k	5.8M
		WMT’06	80–10–10	62k	751k	17.0M	2.0k	61k	3.1k	90k
		WT-103	provided	268k	1.8M	103.2M	3.8k	217k	4.4k	246k
Czech	Indo-European	EuroParl	80–10–10	64k	517k	12.4M	65k	1.6M	65k	1.6M
French	Indo-European	EuroParl	80–10–10	64k	1.6M	51.3M	201k	6.4M	201k	6.3M
German	Indo-European	EuroParl	80–10–10	64k	1.5M	42.3M	192k	5.4M	192k	5.2M
Spanish	Indo-European	EuroParl	80–10–10	64k	1.6M	47.2M	197k	6.0M	197k	5.9M
Finnish	Uralic	Wiki-40B	provided	128k	256k	59.3M	14.1k	3.9M	14.0k	3.2M
Indonesian	Austronesian	Wiki-40B	provided	128k	156k	45.7M	8.7k	3.1M	8.6k	2.5M
Tagalog	Austronesian	Wiki-40B	provided	128k	26k	4.2M	1.5k	270k	1.4k	220k
Turkish	Turkic	Wiki-40B	provided	128k	143k	38.1M	7.8k	2.5M	7.7k	1.9M
Swahili	Niger-Congo	Wikipedia	80–10–10	128k	406k	6.3M	51k	800k	51k	803k

Table 5: Train, dev, and test splits, as well as vocab size, for the language modeling datasets that we use in this paper. If train-dev-test splits were provided, then we used them. Otherwise, we performed a 80–10–10 train-dev-test split. We found a vocab size of 64k to cover more than 98% of the training set for the Indo-European languages, and a vocab size of 62k allowed us to cover 100% in the training set of English WMT’06. For the remaining languages, which had larger vocabularies, we followed Wiki-40B (Guo et al., 2020) and increased the vocab size to 128k.

Hyperparameters. Table 6 shows the optimized β hyperparameter from a grid-search over $\beta \in \{0.006, 0.008, 0.01, 0.02, 0.03, 0.04, 0.05\}$ for both regularizers on all datasets we use. Notably, the best β for variance ranged from 1×10^{-2} to 5×10^{-2} , and the best β for local consistency ranged from 6×10^{-3} to 2×10^{-2} . For use on a new dataset, we recommend starting with 1×10^{-2} , which we found almost always improved perplexity for both regularizers (on these datasets, at least).

Language	Source	UID Regularizer			
		Variance		Local Consistency	
		Best β	Dev Loss	Best β	Dev Loss
English	EuroParl (full dataset)	2×10^{-2}	4.519	8×10^{-3}	4.529
	EuroParl (2M subset)	2×10^{-2}	6.497	1×10^{-2}	6.497
	EuroParl (4M subset)	2×10^{-2}	5.940	1×10^{-2}	5.948
	EuroParl (8M subset)	2×10^{-2}	5.500	8×10^{-3}	5.511
	EuroParl (12M subset)	2×10^{-2}	5.236	8×10^{-3}	5.230
	EuroParl (16M subset)	5×10^{-2}	5.084	2×10^{-2}	5.089
	EuroParl (24M subset)	4×10^{-2}	4.841	2×10^{-2}	4.843
	EuroParl (32M subset)	1×10^{-2}	4.747	1×10^{-2}	4.742
	WMT’06	3×10^{-2}	4.974	1×10^{-2}	4.991
English	WT-103	1×10^{-2}	4.933	8×10^{-3}	4.939
	Czech	EuroParl	3×10^{-2}	5.388	1×10^{-2}
French	EuroParl	1×10^{-2}	4.161	6×10^{-3}	4.162
German	EuroParl	2×10^{-2}	4.782	8×10^{-3}	4.779
Spanish	EuroParl	3×10^{-2}	4.539	1×10^{-2}	4.550
Finnish	Wiki-40B	1×10^{-2}	5.811	6×10^{-3}	5.819
Indonesian	Wiki-40B	3×10^{-2}	5.808	8×10^{-3}	5.809
Tagalog	Wiki-40B	4×10^{-2}	6.319	8×10^{-3}	6.319
Turkish	Wiki-40B	3×10^{-2}	6.119	8×10^{-3}	6.121
Swahili	Wikipedia	2×10^{-2}	5.555	6×10^{-3}	5.546

Table 6: Best β hyperparameters and dev losses for all experiments.