

# Implicit Geometry of Next-token Prediction: From Language Sparsity Patterns to Model Representations

Yize Zhao<sup>†</sup>, Tina Behnia<sup>†</sup>, Vala Vakilian & Christos Thrampoulidis

Department of Electrical & Computer Engineering

University of British Columbia

Vancouver, Canada

{zhaoyize, tina.behnia, vaalaa, cthrampo}@ece.ubc.ca

## Abstract

Next-token prediction (NTP) over large text corpora has become the go-to paradigm to train large language models. Yet, it remains unclear how NTP influences the mapping of linguistic patterns to geometric properties of the resulting model representations. We frame training of large language models as soft-label classification over sparse probabilistic label vectors, coupled with an analytical approximation that allows unrestricted generation of context embeddings. This approach links NTP training to rank-constrained, nuclear-norm regularized optimization in the logit domain, offering a framework for analyzing the geometry of word and context embeddings. In large embedding spaces, we find that NTP implicitly favors learning logits with a sparse plus low-rank structure. While the sparse component captures the co-occurrence frequency of context-word pairs, the orthogonal low-rank component, which becomes dominant as training progresses, depends solely on the sparsity pattern of the co-occurrence matrix. Consequently, when projected onto an appropriate subspace, representations of contexts that are followed by the same set of next-tokens collapse—a phenomenon we term subspace-collapse. We validate our findings on synthetic and small-scale real language datasets. Finally, we outline potential research directions aimed at deepening the understanding of NTP’s influence on the learning of linguistic patterns and regularities.

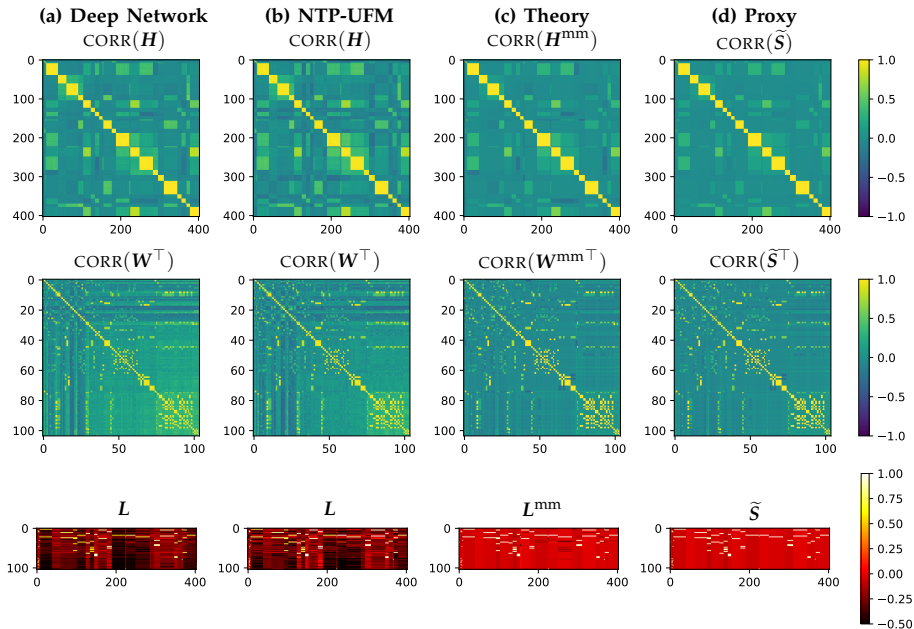
## 1 Introduction

Next-token prediction (NTP) is the preferred training paradigm for state-of-the-art language models. The process, elegantly simple, uses a large training corpus to minimize, for each context  $\mathbf{z}_{<t} \in \mathcal{V}^{t-1}$  of  $t - 1$  preceding tokens, the cross-entropy (CE) loss between the model’s predicted conditional probability distribution over potential next tokens from a vocabulary  $\mathcal{V}$  and the one-hot encoded actual next token  $z_t \in \mathcal{V}$ . The model’s conditional distribution is defined through a softmax map applied to logits  $\ell_{<t}(\mathbf{W}, \boldsymbol{\theta}) = \mathbf{W}\mathbf{h}_{\boldsymbol{\theta}}(\mathbf{z}_{<t})$ , which are generated by mapping *context embeddings*  $\mathbf{h}_{\boldsymbol{\theta}}(\mathbf{z}_{<t}) \in \mathbb{R}^d$ —the neural network’s  $d$ -dimensional representations of contexts  $\mathbf{z}_{<t}$ —using a matrix  $\mathbf{W} \in \mathbb{R}^{|\mathcal{V}| \times d}$  of *word embeddings*.

Rooted in the foundational works of Shannon (1948) and inspired by the “co-occurrence statistics” and “distributional hypothesis” (Harris, 1954), the NTP paradigm suggests that a word’s meaning is defined by its context. This principle underlies both classical vector-space models (Schütze et al., 2008) and neural language models (Turian et al., 2010; Bengio et al., 2000; Baroni et al., 2014; Bengio & Bengio, 2000; Mikolov et al., 2013b), propelling the development of today’s sophisticated large language models (Radford et al., 2018; 2019; Brown et al., 2020). However, a fundamental question remains:

*How does the NTP learning objective shape the relationship between language statistics and the geometry of model representations?*

<sup>†</sup> Equal contribution.



**Figure 1: Can we predict the geometry of a model’s context and word representations based on the structure of the training data? This figure demonstrates that it is possible.** Panel (a) shows the cosine similarity of context and word embeddings,  $\text{CORR}(H)$  and  $\text{CORR}(W^T)$ , at the end of training a four-layer transformer on the Simplified TinyStories dataset, long enough for the NTP loss to converge to its empirical entropy lower-bound. Panel (b) illustrates  $\text{CORR}(H)$  and  $\text{CORR}(W^T)$  found by training our analysis model, NTP-UFM (see Eq. (1)), on the same data. This model abstracts the neural network and allows to predict the geometry analytically. Panel (c) shows this analytical prediction of the geometry that is obtained purely from the data. Specifically,  $H^{\text{mm}}$  and  $W^{\text{mm}}$  are determined by the right/left singular factors (Claim (C2)) of the low-rank (max-margin) component  $L^{\text{mm}}$  of logits (Claim (C1)). Finally, Panel (d) provides an easy to compute heuristic proxy for the embeddings’ geometry based on computing overlaps between the support sets of context’s next-words. Details in Secs. 1.2 and 5.

This inquiry, which we term the *implicit geometry* of NTP—so named because the NTP objective does not explicitly impose any such relationship—explores how distances and angles within the neural network’s  $d$ -dimensional representational space correlate with linguistic patterns at the end of training.

We postulate that understanding this implicit geometry is essential to shedding light on functional principles of large language models, since NTP is used for training across diverse architectures, from LSTMs to transformers and state-space models. Specifically, exploring how optimization under NTP shapes representations of words and contexts, which empirically mirror complex human-like patterns, not only has scientific merit but could also enhance model interpretability and explainability. This is crucial for understanding how biases in training data influence learned representations and how such biases can be mitigated. At the same time, revealing how implicit geometry correlates with language statistics could help us reverse-engineer optimization algorithms and the training loss itself to achieve more desirable structures of model representations, addressing challenges such as statistical imbalances in language data. Finally, understanding how state-of-the-art models internalize language to form representations might also enhance our grasp of language itself.

Answering the question of implicit geometry is inherently challenging, as representations emerge from training complex models (in terms of architecture and size) on diverse, large-scale datasets (varying in source, size, and tokenization) with differing optimization hyperparameters (e.g., learning rate, weight decay, number of iterations). Nonetheless, we demonstrate that characterizing the implicit geometry—at least at a “macroscopic level”—is feasible in a doubly-asymptotic regime of sufficiently large models and extended training periods.

Specifically, this paper develops an analytical framework to characterize the implicit geometry induced by NTP training on language datasets. Drawing inspiration from seminal studies on the geometry of deep model representations in image recognition (Papayan et al.,

2020), our framework distinguishes itself from previous studies on language representations by not concentrating on specific architectures such as transformers. Instead, we assume that the model has adequate representation capacity and undergoes effective optimization, making it possible to minimize the NTP loss to its entropy lower-bound. This approach isolates the influence of NTP itself—rather than architectural nuances—in shaping the implicit geometry of the language model.

The framework reveals the key role that sparsity patterns in language statistics (Thrampoulidis, 2024) play in shaping the implicit geometry. Building on and extending recent work (Thrampoulidis, 2024), we demonstrate that the recurrence of only a few tokens from the entire vocabulary as next-tokens in specific contexts leads to an implicit bias in NTP training. This bias favors a matrix of logits that develops a *sparse plus low-rank* structure during training. The sparse component of this matrix captures the probabilities of co-occurring words and contexts, while the dominant low-rank component reflects the sparsity pattern of the co-occurrence matrix. Consequently, embeddings of contexts followed by similar sets of words—regardless of their frequencies—increasingly align during training. Simultaneously, the sparse component of the logits ensures accurate prediction of the correct frequencies for each context, which is crucial for achieving the entropy lower bound. Importantly, both the sparse and low-rank components can be computed purely from the input data, enabling us to predict the geometric structure of a model’s representations trained on the same data a priori, without the need for training; see Fig. 1 for an illustration. Overall, our framework introduces a novel perspective, distinct from traditional analyses of word representations like those in the Word2Vec model (Levy & Goldberg, 2014).

## 1.1 Methodology

Our methodology integrates three foundational modeling concepts as follows:

Firstly, following (Thrampoulidis, 2024) we frame NTP as soft-label classification with CE loss applied to *sparse* probabilistic label vectors. This isolates  $m$  *distinct* contexts, which could be repeated multiple times throughout the training corpus, and assigns to each a sparse  $V = |\mathcal{V}|$ -dimensional conditional-probability label vector  $\hat{p}_j$ , reflecting the frequency of each token following context  $j \in [m]$ . The sparsity of  $\hat{p}_j$  indicates that certain tokens, which we refer to as *off-support* tokens for the specific  $j$ -th context, never follow this context.

The second concept facilitates a tractable analysis of context embeddings by assuming expressive (enough) neural networks can produce unconstrained embeddings  $h_j \in \mathbb{R}^d$ , independent of the architecture’s specific complexities (Yang et al., 2017; Mixon et al., 2020). This redefines NTP as a minimization, over word and context embedding matrices  $W$  and  $H$ , of the NTP loss across a training set  $\mathcal{T}$  determined by the *sparse* matrix of next-token conditional probabilities

$$P = P(\mathcal{T}) = [\hat{p}_1, \hat{p}_2, \dots, \hat{p}_m] \in [0, 1]^{V \times m}.$$

Overall, this leads to the following log-bilinear model:

$$\min_{W \in \mathbb{R}^{V \times d}, H \in \mathbb{R}^{d \times m}} \mathcal{L}_{\text{NTP}}(WH; P). \tag{1}$$

This way, our goal to study word-word, context-context, and word-context geometric relationships becomes that of characterizing the Gram matrices  $G_W = WW^\top$  and  $G_H = H^\top H$ , as well as the logit matrix  $L = WH$  at the minimizers of Eq. (1). This task is complicated by the non-convex nature of the minimization and the sparsity of the probabilistic label vectors. Specifically, as we show, the sparsity of  $P$  may lead to multiple minimizers, potentially making geometric characterization ambiguous.

To address this, we leverage a third concept: focusing on specific minimizers identified through the *regularization path* (Rosset et al., 2003). This entails following the solution trajectory of the empirical risk minimization in Eq. (1) as an additive ridge regularization for  $W$  and  $H$  diminishes to zero. For the purpose of comparing the analysis outcomes to our numerical evaluations, we interpret the prediction obtained from the regularization path analysis as a proxy for the solution found by gradient-based optimization when the NTP loss approaches its empirical entropy lower bound.

## 1.2 Summary of findings

Building on the above methodology, our analytical framework leads to the following results:

**Formulation in logit space.** Sec. 3 presents an equivalent formulation of the NTP objective

$$\mathcal{L}_{\text{NTP}}(\mathbf{W}\mathbf{H}; \mathbf{P}) + \lambda \|\mathbf{W}\|^2 + \lambda \|\mathbf{H}\|^2$$

into the logit space, given in terms of a rank-constrained and nuclear-norm regularized minimization. From this, word and context matrices can be obtained through matrix factorization of logits  $\mathbf{L}$ .

**Logit Convergence.** Focusing on  $d \geq V^1$ , in Sec. 4, we demonstrate that the logit matrix  $\mathbf{L}_\lambda$  for regularization  $\lambda \rightarrow 0$  (which is proxy for training iteration  $k \rightarrow \infty$ ), behaves for some  $R(\lambda) \rightarrow \infty$  as shown by the following claim (see Fig. 2 for visualization):

**(C1) Logits' sparse plus low-rank decomposition:** As  $\lambda \rightarrow 0$ , for some  $R(\lambda) \rightarrow \infty$ :

$$\mathbf{L}_\lambda \approx \underbrace{\mathbf{L}^{\text{in}}}_{\substack{\text{sparse component:} \\ \text{sets in-support token probabilities}}} + R(\lambda) \cdot \underbrace{\mathbf{L}^{\text{mm}}}_{\substack{\text{low-rank component:} \\ \text{separates in- from off-support tokens}}}, \quad (2)$$

where the orthogonal components  $\mathbf{L}^{\text{in}}$  and  $\mathbf{L}^{\text{mm}}$  have distinct operational roles.

$\mathbf{L}^{\text{in}}$  inherits the sparsity of the data matrix  $\mathbf{P}$  and encodes information regarding frequencies of in-support tokens. In contrast,  $\mathbf{L}^{\text{mm}}$  is independent of these frequencies and is influenced only by the sparsity pattern  $\mathbf{S} \in \{0, 1\}^{V \times m}$ . Thus, its value is guided by both the observed ‘company’ of each context (in-support tokens), but not their frequency, and, by the ‘company’ it lacks (off-support tokens). More specifically,  $\mathbf{L}^{\text{mm}}$  functions to maximize the logit margin between in-support and off-support tokens, where the margin is defined in terms of the nuclear norm that promotes low rank, rather than the Euclidean norm. To highlight this property, we refer to  $\mathbf{L}^{\text{mm}}$  as the ‘low-rank’ component, though it could also be described as the ‘max-margin’ component. The decomposition in Eq. (2) also shows the logits become unbounded in norm when projected on the subspace of  $\mathbf{L}^{\text{mm}}$ .

**Context and word embeddings’ geometry.** Similar to logits, word/context embeddings  $\mathbf{W}$  and  $\mathbf{H}$  also grow in norm as  $\lambda \rightarrow 0$ . This occurs in a way that simultaneously guarantees the resulting logit matrix abides by Claim (C1), which leads to the following additional claims:


**(C2) Directional convergence:** Word and context embedding matrices converge *directionally* to matrices  $\mathbf{W}^{\text{mm}} := \mathbf{U}\mathbf{\Sigma}^{1/2}\mathbf{R}$ ,  $\mathbf{H}^{\text{mm}} := \mathbf{R}^\top \mathbf{\Sigma}^{1/2}\mathbf{V}^\top$ , where  $\mathbf{R}$  is a rotation matrix and  $\mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top$  is the singular value decomposition (SVD) of the low-rank max-margin component  $\mathbf{L}^{\text{mm}}$  in Eq. (2). Concretely, letting  $\bar{\mathbf{A}} := \mathbf{A}/\|\mathbf{A}\|$  for any matrix  $\mathbf{A}$ :

$$\bar{\mathbf{G}}_{\mathbf{W}} \rightarrow \mathbf{U}\bar{\mathbf{\Sigma}}\mathbf{U}^\top \quad \text{and} \quad \bar{\mathbf{G}}_{\mathbf{H}} \rightarrow \mathbf{V}\bar{\mathbf{\Sigma}}\mathbf{V}^\top. \quad (3)$$

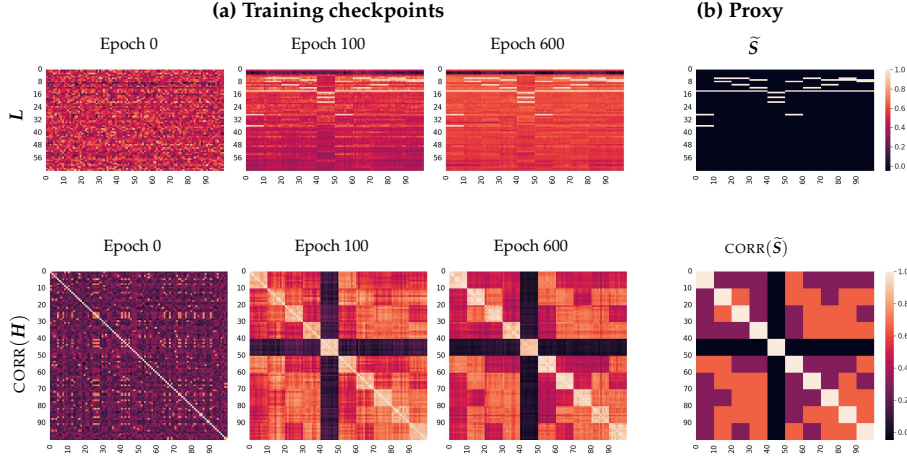
**(C3) Subspace collapse:** Since  $\mathbf{L}^{\text{mm}}$  depends only on the sparsity patterns of next-token distributions (not their frequencies), a consequence of (C2) is a property termed *subspace collapse*: embeddings  $\mathbf{h}_j, \mathbf{h}_{j'}$  of contexts  $j \neq j' \in [m]$  that are followed by the same set of next-tokens (although their frequencies may differ), converge to the same limiting direction, i.e., for all  $j, j' \in [m]$

$$\text{support}(\hat{\mathbf{p}}_j) = \text{support}(\hat{\mathbf{p}}_{j'}) \implies \cos(\mathbf{h}_j, \mathbf{h}_{j'}) = \mathbf{h}_j^\top \mathbf{h}_{j'} / (\|\mathbf{h}_j\| \|\mathbf{h}_{j'}\|) \rightarrow 1. \quad (4)$$

<sup>1</sup>This does *not* constrain  $d$  relative to the number of distinct contexts  $m$ , which could be significantly larger than the vocabulary size  $V$ . Also, despite large embedding dimensions, our results still reveal how the geometry of learned word embeddings encodes fine-grained similarities between words as materialized in the patterns of language data; e.g., refer to Fig. 8 for a visualization.



**Figure 2:** In the limit of vanishing regularization, the ridge-regularized NTP objective solution  $\mathbf{L}_\lambda$  (aka the logit matrix) decomposes into two orthogonal components. The component  $\mathbf{L}^{\text{in}}$  inherits the sparsity of the matrix  $\mathbf{P}$  of next-token conditional probabilities. In the orthogonal complement,  $\mathbf{L}_\lambda$  diverges in norm but converges in direction to  $\mathbf{L}^{\text{mm}}$ , which results from a low-rank-promoting nuclear-norm minimization that depends only on the sparsity pattern of  $\mathbf{P}$ .



**Figure 3:** Similar to Fig. 1, this time on a 12-layer TF trained on a subset of 100 stories from the TinyStories dataset. Here, computing the theoretical prediction  $H^{\text{mm}}$  is computationally expensive. Thus, we compare the embeddings geometry with the Proxy (P). Details in Sec. 5.

**(C4) Soft-label interpolation:** When projected on the subspace of in-support tokens, the logits  $WH$  interpolate the corresponding soft-labels ensuring that the NTP loss reaches the entropy lower bound. Concretely, for all  $j \in [m]$ ,

$$(w_z - w_{z'})^\top h_j = \log \left( \hat{p}_{j,z} / \hat{p}_{j,z'} \right), \quad \forall z, z' \in \text{support}(\hat{p}_j). \quad (5)$$

Since the subspace of in-support tokens is orthogonal to the subspace of  $L^{\text{mm}}$ , the subspace collapse (C3) does *not* prevent logits  $w_z^\top h_j$  and  $w_{z'}^\top h_j$  of in-support tokens  $z \in \text{support}(\hat{p}_j) = \text{support}(\hat{p}_{j'})$  to interpolate (in the sense of Eq. (5)) potentially different conditional probabilities (soft-labels)  $\hat{p}_{j,z}$  and  $\hat{p}_{j',z'}$ , respectively.

**Experiments.** Sec. 5 presents experiments on controlled settings that validate our findings. An example is shown in Fig. 1: The experiment involves training a 4-layer transformer (TF) on a training set extracted and curated from the TinyStories dataset (Eldan & Li, 2023), with the following characteristics: a vocabulary size  $V = 104$  and a total number of contexts  $n \approx 3050$ , out of which  $m \approx 400$  are distinct. We choose TF embedding dimension  $d = 128 > V$  and train until the empirical entropy lower bound is reached within an order of  $10^{-4}$ . The leftmost Panel (a) depicts the geometry of the context embeddings  $H$  (Top), word embeddings  $W$  (Middle), and logits  $L$  (Bottom) learned by the TF at the end of training. Panel (b) compares these with the same quantities learned by the log-bilinear model in Eq. (1) evaluated on the same training dataset. The apparent resemblance of the patterns validates that the proposed log-bilinear analysis model is a good proxy for the TF model. Panel (c) compares the TF output to our analytical predictions: Comparing TF’s context/word embeddings and logits to  $H^{\text{mm}}/W^{\text{mm}}$  and  $L^{\text{mm}}$  confirms Claims (C2) and (C1). The bright yellow regions in the Gram matrices of embeddings help visualize the subspace collapse Claim (C3). See Sec. 5 for additional details and verification of Claim (C4).

Finding the analytical prediction (C2), illustrated in Panel (c), requires solving a semi-definite program for  $L^{\text{mm}}$  and computing its SVD. In large scales, this can be computationally prohibitive, motivating the following heuristic proxy for the embedding geometry.

**(P) Empirical proxy for directional convergence:** Let  $\tilde{S} := (\mathbb{I}_V - \frac{1}{V}\mathbb{1}_V\mathbb{1}_V^\top)S$  be the (column-wise) centered support-set matrix. The directional component of word and context embedding matrices can be well-approximated by  $\tilde{S}$  as follows:

$$W^{\text{mm}}W^{\text{mm}\top} \approx \tilde{S}\tilde{S}^\top \quad \text{and} \quad H^{\text{mm}\top}H^{\text{mm}} \approx \tilde{S}^\top\tilde{S}. \quad (6)$$

Panel (d) of Fig. 1 suggests proxy (P) as a good enough approximator of the embeddings structure that simply only depends on the sparsity pattern  $S$  of the next-tokens’ conditional probability matrix.



A second example supporting proxy **(P)** is shown in Fig. 3: Here, we train a larger TF on a subset of the *TinyStories*, where computing  $L^{\text{mm}}$  and its SVD is expensive. Instead, we observe that the structure of the support set closely captures the structure of the learned embeddings: Along the diagonal blocks, where contexts with the exact same support sets are situated, the context embeddings align closely (subspace collapse in Claim **(C3)**). Conversely, when support sets have zero intersection (dark entries in Panel (b)), context embeddings exhibit low correlation. This finding is consistent with classical intuitions of the “distributional assumption” of words/contexts and is formalized by our framework via the implicit bias of NTP to promote low-rank logits subject to margin conditions for in/off-support tokens.

## 2 Related work

We pinpoint three main related research areas; see App. A for more in-depth discussion.

First, our research conceptually mirrors the seminal work by [Levy & Goldberg \(2014\)](#) who framed the Skip-Gram with Negative Sampling (SGNS) training objective of Word2Vec ([Mikolov et al., 2013b;a](#)) as weighted matrix factorization. Specifically for large  $d$ , they demonstrated that SGNS implicitly factorizes the pointwise mutual information (PMI) matrix. Their analysis relies on the fact that Word2Vec architecture already is a log-bilinear model, while the log-bilinear model in (1) is only used by us as an analytical proxy for more complex architectures. Also, different to them, we focus on NTP, which employs softmax instead of sigmoids as in SGNS. (However, our analysis applies also to SGNS; see App. A.4). More importantly, we contribute a fresh perspective on this line of inquiry by: (i) confronting the sparsity of probabilistic labels head-on, identifying that it leads to diverging embeddings, a scenario where setting the loss gradient to zero, as in [Levy & Goldberg \(2014\)](#), is infeasible; (ii) examining embeddings through the lens of the regularization path—a surrogate for gradient descent optimization— which unveils that embeddings emerge from the factorization of  $L^{\text{in}} + RL^{\text{mm}}$ , where  $L^{\text{in}}$  captures frequencies akin to the PMI, while the directional component  $L^{\text{mm}}$ , becoming dominant as weights diverge, reflects the explicit sparsity patterns of the context-word co-occurrences. We envision this fresh perspective could similarly motivate further research leveraging the geometric insights of embeddings to uncover linguistic phenomena such as the linear relationships underlying word analogies ([Mikolov et al., 2013a](#); [Pennington et al., 2014](#)).

Second, our exploration of NTP’s regularization path is inspired by the *implicit bias/regularization* research on the preferred solutions of optimizers like GD in overparameterized systems. [Ji & Telgarsky \(2020\)](#) showed that GD’s trajectory in linear one-hot encoding models aligns with the regularization path, providing a lens for examining GD dynamics. [Thrampoulidis \(2024\)](#) recently extended this analysis to NTP, framing it as sparse soft-label classification. By lifting their assumption of fixed context embeddings, we delve into the more complex non-convex domain, establishing explicit connections between the implicit geometry and linguistic patterns, which we also verify empirically.

Finally, our research intersects with the study of *neural-collapse (NC)*, which delves into the geometry of last-layer features and weights in deep networks trained in the interpolating regime ([Papayan et al., 2020](#)), utilizing the unconstrained features model (UFM) for analysis ([Mixon et al., 2020](#); [Fang et al., 2021](#); [Zhu et al., 2021](#)). Our work is particularly aligned with [Thrampoulidis et al. \(2022\)](#): We also examine the UFM’s regularization path but within the NTP framework, extending their one-hot classification findings, which can be seen as special cases of ours. Additionally, our findings have some parallels with the multilabel NC geometry explored in [Li et al. \(2023\)](#); [Fisher et al. \(2024\)](#). However, our setting is more general raising stringent assumptions on the label distribution. For detailed comparison, see App. A.

## 3 Formulation

**Notations.** Throughout, lowercase and uppercase bold letters (e.g.,  $a$  and  $A$ ) represent vectors and matrices, respectively. We use  $\bar{a}$  and  $\bar{A}$  to denote vectors/matrices normalized by their Euclidean norm. We denote  $A[i, j]$  the  $(i, j)$ -th entry of matrix  $A$  and  $a_j$  its  $j$ -

th column. We let  $\mathcal{R}(A)$  and  $\mathcal{N}(A)$  its range-space (aka column space) and null-space, respectively.  $\langle \cdot, \cdot \rangle$  and  $\|\cdot\|$  denote Euclidean inner product and norm, respectively. We use  $\|\cdot\|_*$  to denote the nuclear-norm (i.e. sum of singular values).  $\mathbb{I}_V$  represents the identity matrix of size  $V$  and  $\mathbb{1}_V$ , the all ones vector of size  $V \times 1$  (subscripts are removed when clear from context).  $\Delta^{V-1}$  denotes the  $V$ -dimensional unit simplex and  $S(\cdot) : \mathbb{R}^V \rightarrow \Delta^{V-1}$  the softmax map:

$$S(\mathbf{a}) = [S_1(\mathbf{a}), \dots, S_V(\mathbf{a})]^\top, \quad \text{with } S_v(\mathbf{a}) = \frac{\exp(\mathbf{e}_v^\top \mathbf{a})}{\sum_{v' \in [V]} \exp(\mathbf{e}_{v'}^\top \mathbf{a})},$$

where  $\mathbf{e}_v$  is the  $v$ -th standard basis vector in  $\mathbb{R}^V$ . We also denote  $\tilde{\mathbf{e}}_j$  the  $j$ -th standard basis vector in  $\mathbb{R}^m$ . All logarithms are natural logarithms (base  $e$ ).

### 3.1 NTP objective as sparse soft-label classification

We let  $\mathcal{V} = [V] := \{1, \dots, V\}$  represent a finite vocabulary of tokens (we use the terms ‘word’ and ‘token’ interchangeably). We denote by  $\mathbf{z}_{1:t} = (z_1, \dots, z_t)$  a sequence of  $t$  tokens  $z_t \in \mathcal{V}$  and focus, for simplicity, on prediction of the last  $T$ -th token  $z := z_T$  given context  $\mathbf{x} := \mathbf{z}_{1:T-1}$ . For this, we assume access to a training set consisting of  $n$  sequences  $\mathcal{T}_n := \{(x_i, z_i)\}_{i \in [n]}$ , such that  $x_i \in \mathcal{X} := \mathcal{V}^{T-1}$  and  $z_i \in \mathcal{V}$ . This is used to train model  $f_{\theta'} : \mathcal{X} \rightarrow \mathcal{V}$ ,  $f_{\theta'}(\mathbf{x}) = \mathbf{W}h_{\theta}(\mathbf{x})$  parameterized by  $\theta' = \{\mathbf{W}, \theta\}$ , where  $\mathbf{W} \in \mathbb{R}^{V \times d}$  is a decoding matrix and  $\theta$  parameterizes a map  $h_{\theta} : \mathcal{X} \rightarrow \mathbb{R}^d$  from contexts to  $d$ -dimensional embeddings. We impose no restrictions on the specific form of the embedding map, which may, for example, be produced by an MLP, an LSTM, or a TF. We refer to row  $w_v, v \in \mathcal{V}$  of  $\mathbf{W}$  as **word embedding** of token  $v$  and  $h_{\theta}(\mathbf{x})$  as **context embedding** of context  $\mathbf{x}$ . The model is found by minimizing the empirical CE loss  $\text{CE}(\theta') = \frac{1}{n} \sum_{i \in [n]} -\log(S_{z_i}(f_{\theta'}(x_i)))$ .

Following [Thrapoulidis \(2024\)](#), we reframe the NTP training objective as classification over  $m \leq n$  distinct contexts, each associated with a *sparse* probabilistic label vector  $\hat{\mathbf{p}}_j \in \Delta^{V-1}$ . Concretely, we denote  $\bar{x}_1, \dots, \bar{x}_m$  the  $m \leq n$  distinct contexts among the (large number of) total  $n$  contexts  $\mathcal{T}_n$ . Also, we let  $\hat{\pi}_j = \frac{1}{n} \sum_{i \in [n]} \mathbb{1}[x_i = \bar{x}_j]$  denote the empirical probability of distinct context  $\bar{x}_j$ . Accordingly, let  $\hat{\mathbf{p}}_j \in \Delta^{V-1}$  denote the probability vector of conditional next-token distribution, i.e., for all  $z \in \mathcal{V}$ :  $\hat{p}_{j,z} := \frac{\sum_{i \in [n]: x_i = \bar{x}_j} \mathbb{1}[z_i = z]}{\sum_{i \in [n]} \mathbb{1}[x_i = \bar{x}_j]}$ ,  $j \in [m]$ . In words,  $n \cdot \hat{\pi}_j \cdot \hat{p}_{j,z}$  is the number of occurrences of token  $z$  as a follow-up to context  $\bar{x}_j$ . Define the support sets of these probability vectors as  $\mathcal{S}_j := \{z \in \mathcal{V} \mid \hat{p}_{j,z} > 0\}$  and let  $S_j := |\mathcal{S}_j|$ .

It is also convenient to define the **conditional probability matrix**

$$\mathbf{P} = [\hat{\mathbf{p}}_1, \dots, \hat{\mathbf{p}}_m] \in \mathbb{R}^{V \times m},$$

and its corresponding **support matrix**  $\mathbf{S} \in \{0, 1\}^{V \times m}$ , such that  $S[z, j] = 1$ , iff  $z \in \mathcal{S}_j$ . In typical scenarios,  $S_j < V$ ; thus,  $\mathbf{P}$  and  $\mathbf{S}$  are *sparse* matrices. For given context index  $j \in [m]$ , we say word  $z$  is **in-support token** if  $S[z, j] = 1$  (eqv.  $z \in \mathcal{S}_j$ ); otherwise, we say  $z$  is **off-support token**.

With the above notation, the training loss becomes ([Thrapoulidis, 2024](#)),

$$\text{CE}(\theta') = - \sum_{j \in [m]} \hat{\pi}_j \sum_{z \in \mathcal{V}} \hat{p}_{j,z} \log(S_z(\mathbf{W}h_{\theta}(\bar{x}_j))). \quad (7)$$

This is lower bounded by the empirical  $T$ -gram entropy (referred to hereafter as entropy) of the data ([Shannon, 1948](#)), i.e., for all  $\theta'$ :

$$\text{CE}(\theta') \geq \mathcal{H} := - \sum_{j \in [m]} \hat{\pi}_j \sum_{z \in \mathcal{V}} \hat{p}_{j,z} \log(\hat{p}_{j,z}).$$

### 3.2 Unconstrained features model for NTP training

To gain insights into the geometry of solutions to CE minimization in Eq. (7), we assume sufficient model expressivity, allowing us to optimize embeddings freely, instead of abiding by their architecture-specific parameterization. We formalize this concept below.

**Definition 1** (NTP-UFM). *The unconstrained features model (UFM) for NTP training over training set  $\mathcal{T}_m = \{\hat{\tau}_j, \hat{p}_j\}_{j \in [m]}$  refers to the following log-bilinear optimization problem:*

$$\min_{\mathbf{W}, \mathbf{H}} \text{CE}(\mathbf{WH}) + \frac{\lambda}{2} \|\mathbf{W}\|^2 + \frac{\lambda}{2} \|\mathbf{H}\|^2, \quad (\text{NTP-UFM})$$

where minimization is over word and context embedding matrices  $\mathbf{W} \in \mathbb{R}^{V \times d}$  and  $\mathbf{H} := [\mathbf{h}_1, \dots, \mathbf{h}_m] \in \mathbb{R}^{d \times m}$ , and  $\text{CE}(\mathbf{WH}) := -\sum_{j \in [m]} \hat{\tau}_j \sum_{z \in \mathcal{S}_j} \hat{p}_{j,z} \log(\mathcal{S}_z(\mathbf{W}\mathbf{h}_j))$ .

The key difference of **NTP-UFM** compared to Eq. (7) is that embeddings  $\mathbf{h}_\theta(\bar{x}_j)$  are now optimized unconstrainedly through free variables  $\mathbf{h}_j$ . Furthermore, we have introduced ridge regularization.<sup>2</sup> Although our focus is on understanding the geometry arising from minimizing the unregularized CE objective in Eq. (7), we utilize ridge regularization as a proxy to examine the behavior of gradient descent.<sup>3</sup> Specifically, we use the vanishing-regularization solution of **NTP-UFM** as a proxy for the parameters  $(\mathbf{W}_k, \mathbf{H}_k)$  learned in large iteration  $k \rightarrow \infty$  through GD on the unregularized objective (7).

**Definition 2** (Regularization path). *The regularization path of the **NTP-UFM** minimization corresponds to the solution  $(\mathbf{W}_\lambda, \mathbf{H}_\lambda)$  of the respective log-bilinear program in the limit of vanishing regularization  $\lambda \rightarrow 0$ .*

### 3.3 Reformulation in terms of logits

We now introduce a reformulation of **NTP-UFM** utilizing the logit matrix  $\mathbf{L} = \mathbf{WH}$ . This is grounded on the well-known fact concerning the nuclear norm of a matrix (Srebro et al., 2004; Fazel, 2002):  $\|\mathbf{L}\|_* = \min_{\mathbf{L}=\mathbf{WH}} \frac{1}{2} \|\mathbf{W}\|^2 + \frac{1}{2} \|\mathbf{H}\|^2$ .

**Lemma 1.** *Denote logit matrix  $\mathbf{L} = [\ell_1, \dots, \ell_m] \in \mathbb{R}^{V \times m}$  and let  $\mathbf{L}_\lambda$  be a minimizer of*

$$\min_{\mathbf{L}: \text{rank}(\mathbf{L}) \leq d} \left\{ -\sum_{j \in [m]} \hat{\tau}_j \sum_{z \in \mathcal{S}_j} \hat{p}_{j,z} \log(\mathcal{S}_z(\ell_j)) + \lambda \|\mathbf{L}\|_* \right\}, \quad (8)$$

with SVD  $\mathbf{L}_\lambda = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top$ , where  $\mathbf{U} \in \mathbb{R}^{V \times r}$ ,  $\mathbf{\Sigma} \in \mathbb{R}^{r \times r}$ ,  $\mathbf{V} \in \mathbb{R}^{m \times r}$  and  $r = \text{rank}(\mathbf{L}_\lambda) \leq d$ . Then:

- (i) *The optimal cost of **NTP-UFM** is the same as that of Eq. (8).*
- (ii)  *$(\mathbf{W}_\lambda, \mathbf{H}_\lambda)$  minimizes **NTP-UFM** if and only if there exists a minimizer  $\mathbf{L}_\lambda$  of Eq. (8) such that:*

$$\mathbf{W}_\lambda \mathbf{H}_\lambda^\top = \mathbf{L}_\lambda, \quad \mathbf{W}_\lambda \mathbf{W}_\lambda^\top = \mathbf{U}\mathbf{\Sigma}\mathbf{U}^\top, \quad \text{and} \quad \mathbf{H}_\lambda^\top \mathbf{H}_\lambda = \mathbf{V}\mathbf{\Sigma}\mathbf{V}^\top.$$

The constraint  $\text{rank}(\mathbf{L}) \leq d$  in Eq. (8) ensures the logit matrix can be factorized as  $\mathbf{WH}$  with inner factor dimension being  $d$ . Thus, in general, Eq. (8) is non-convex.

## 4 Analysis of unconstrained features model for NTP training

This section analyzes the regularization path of Eq. (8) and, subsequently, of **NTP-UFM**. Throughout, we assume  $d \geq V$ ,<sup>4</sup> in which case the constraint  $\text{rank}(\mathbf{L}) \leq d$  becomes redundant. All proofs are deferred to the appendix and numerical evaluations to Sec. 4.5.

### 4.1 NTP-SVM logits

The following nuclear norm-minimization problem plays a key role in our results.

<sup>2</sup>For simplicity, we apply equal regularization to  $\mathbf{W}$  and  $\mathbf{H}$ . While our results can easily extend to the general scenario, such extensions do not yield additional insights.

<sup>3</sup>The equivalence between GD and regularization paths has been rigorously established in certain convex settings (Ji et al., 2020), but it may not hold generally. Our setting is non-convex, so we use this equivalence as a heuristic proxy; see future work discussion in Sec. 6.

<sup>4</sup>This can be slightly improved to  $d \geq V - 1$ : Strong convexity of ridge-regularization and translational invariance of softmax imply any minimizer  $\widehat{\mathbf{W}}$  of **NTP-UFM** is centered, i.e.,  $\mathbf{1}_V^\top \widehat{\mathbf{W}}_\lambda = 0$ . This implies  $\mathbf{1}_V^\top \mathbf{L} = 0$  can be added in Eq. (8) without changing Lemma 1.



**Definition 3.** Define the NTP-SVM logit matrix  $L^{\text{mm}}$  as solution to the following optimization:

$$\begin{aligned} L^{\text{mm}} \in \arg \min_{L \in \mathbb{R}^{V \times m}} \|L\|_* & \quad (\text{NTP-SVM}_\star) \\ \text{subj. to } L[z, j] - L[z', j] = 0, \quad \forall j \in [m], z \neq z' \in \mathcal{S}_j, & \quad (9) \\ L[z, j] - L[v, j] \geq 1, \quad \forall j \in [m], z \in \mathcal{S}_j, v \notin \mathcal{S}_j. & \quad (10) \end{aligned}$$

The optimization in  $\text{NTP-SVM}_\star$  is a semidefinite program (SDP) and always feasible, as shown by the straightforward feasibility of the centered support matrix  $\tilde{S} = (\mathbb{I}_V - \frac{1}{V} \mathbf{1}_V \mathbf{1}_V^\top) S$ . Its solution  $L^{\text{mm}}$  depends solely on the support matrix  $S$  of the data. Interestingly, we prove in Sec. 4.4 that  $\tilde{S}$  solves  $\text{NTP-SVM}_\star$  under perfect symmetry assumptions on  $S$ . More generally, our empirical results suggest that  $\tilde{S}$  closely approximates the true solution  $L^{\text{mm}}$ , supporting the use of Proxy (P) as explained in Sec. 4.4.2.

We remark that while the constraints are defined for all pairs of tokens, for each context  $j$ , the number of linearly independent constraints imposed by Eqs. (9) and (10) is  $V - 1$ . This reduction is possible because the constraints can be reformulated relative to an anchor  $z_j \in \mathcal{S}_j$  given the equality constraints (9).

Finally, when  $\mathcal{S}_j = 1, \forall j \in [m]$ , then Eq. (9) vanishes and Eq. (10) lower bounds the margin between the correct token and the rest. This constraint appears in the classical hard-margin SVM, which explains our naming. Instead, the NTP setting results in additional equality constraints and minimizes nuclear-norm rather than Euclidean norm.

## 4.2 Regularization-path analysis of logits

We now characterize the solutions of (8) when regularization vanishes and  $d \geq V$ .

**Theorem 1.** Suppose  $d \geq V$ . In the limit of vanishing regularization, the solution  $L_\lambda$  of optimization (8) diverges in norm and converges in direction to  $L^{\text{mm}}$ . Formally,  $\lim_{\lambda \rightarrow 0} \|L_\lambda\| = +\infty$ , and, there exists minimizer  $L^{\text{mm}}$  of  $\text{NTP-SVM}_\star$  such that

$$\lim_{\lambda \rightarrow 0} \left\| \frac{L_\lambda}{\|L_\lambda\|_*} - \frac{L^{\text{mm}}}{\|L^{\text{mm}}\|_*} \right\| = 0.$$

Thus, with vanishing regularization,  $\text{NTP-SVM}_\star$  captures the structure of the logits. By Lemma 1, it also captures the structure of  $W$  and  $H$ . We explore this further in Sec. 4.3.

Now, recall that  $\text{NTP-SVM}_\star$  only depends on the support matrix of the training data, raising the questions: How do the probabilities  $\hat{p}_{j,z}$  influence the learned logit matrix and what implications this has for the loss? The theorem below complements Thm. 1 and addresses these questions. To state the result, define the following matrix subspace, as the span of rank-one matrices each being zero except entries  $(z, j)$  and  $(z', j)$ , which equal 1 and  $-1$  respectively for all context indices  $j \in [m]$  and in-support word indices  $z, z' \in \mathcal{S}_j$ :

$$\mathcal{F} = \mathcal{F}(S) = \text{span} \left( \{ (e_z - e_{z'}) \tilde{e}_j^\top : z \neq z' \in \mathcal{S}_j, j \in [m] \} \right) \subset \mathbb{R}^{V \times m}.$$

Note for any  $L$ , the projection  $\mathcal{P}_{\mathcal{F}}(L)$  onto  $\mathcal{F}$  is sparse with support matching  $S$ .

**Theorem 2.** Under the setting of Theorem 1, the loss approaches its lower bound  $\lim_{\lambda \rightarrow 0} \text{CE}(L_\lambda) = \mathcal{H}$ . Additionally,  $\lim_{\lambda \rightarrow 0} \mathcal{P}_{\mathcal{F}}(L_\lambda) = L^{\text{in}}$ , where  $L^{\text{in}} \in \mathcal{F}$  is the unique  $L \in \mathcal{F}$  that guarantees differences of logits equal log-odds, i.e.,

$$L[z, j] - L[z', j] = \log \left( \frac{\hat{p}_{j,z}}{\hat{p}_{j,z'}} \right), \quad \forall z \neq z' \in \mathcal{S}_j, j \in [m].$$

The loss reaches the lower-bound  $\mathcal{H}$  because on the subspace  $\mathcal{F}$ ,  $L_\lambda$  satisfies the log-odds linear equations of the theorem. This gives rise to Claim (C4): the model outputs the correct probabilities for in-support tokens. There is no contradiction of this claim to Thm. 1 because  $L^{\text{mm}}$  belongs to the orthogonal complement  $\mathcal{F}_\perp$  (which follows from Eq. (9); see App. B.3.1).

Combining the two theorems,  $L_\lambda$  converges to  $L^{\text{in}}$  in  $\mathcal{F}$  and diverges in  $\mathcal{F}_\perp$ , where it converges directionally to  $L^{\text{mm}}$ . This justifies Claim (C1).

### 4.3 Geometry of embeddings

In view of Lemma 1 we can obtain the embeddings by factorizing the logit matrix. Since logits diverge, the same is true for the embeddings, which also converge directionally.

**Corollary 1.** Denote  $(W_\lambda, H_\lambda)$  any minimizer of *NTP-UFM*. Consider the SVD  $L^{\text{mm}} = U\Sigma V^\top$ . Then, using  $\bar{\cdot}$  notation to denote normalized quantities, in the limit of vanishing regularization

$$\overline{W_\lambda W_\lambda^\top} \rightarrow U\bar{\Sigma}U^\top =: \bar{G}_W^{\text{mm}}, \quad \overline{H_\lambda^\top H_\lambda} \rightarrow V\bar{\Sigma}V^\top =: \bar{G}_H^{\text{mm}}, \quad \text{and} \quad \overline{W_\lambda H_\lambda} \rightarrow U\bar{\Sigma}V^\top.$$

This corollary supports Claim (C2). Another direct consequence of it is Claim (C3): embeddings of contexts whose support sets are identical asymptotically collapse to the same embedding as  $\lambda \rightarrow 0$ .

**Proposition 1** (Subspace collapse). Assume *NTP-SVM\** is feasible. There exists minimizer  $H = [h_1, \dots, h_m]$  of *NTP-UFM* such that for all contexts  $j, j' \in [m]$  with same support set  $S_j = S_{j'}$ , their embeddings are same in direction in the limit of vanishing regularization, i.e.,  $\lim_{\lambda \rightarrow 0} \|\bar{h}_j - \bar{h}_{j'}\| = 0$ .

This is consistent with our empirical observations in Figs. 1 and 3: As the loss approaches its lower bound, context embeddings with the same support set converge in the same direction and exhibit maximum correlation. This is indicated by the bright diagonal blocks in  $\text{CORR}(H)$  heatmaps.

It is important to reiterate that Prop. 1 does not prevent the minimizer  $L_\lambda = W_\lambda H_\lambda$  from accurately recovering the distinct soft-label values  $\hat{p}_j, \hat{p}_{j'}$  and achieving the empirical entropy lower-bound  $\mathcal{H}$  as  $\lambda \rightarrow 0$ : The collapse occurs on subspace  $\mathcal{F}_\perp$ , orthogonal to the data subspace  $\mathcal{F}$  on which the logit matrix satisfies the log-odds constraints indicated in Thm. 2; see Secs. 4.5 and 5 for numerical verifications.

### 4.4 On the solution of *NTP-SVM\**

To better understand how the SVD factorization informs the geometry, we investigate further how the structure of  $L^{\text{mm}}$  (solution to *NTP-SVM\**) depends on the support matrix  $S$ .

We start in Sec. 4.4.1 with an idealized, perfectly symmetric setting where each of the  $\binom{V}{k}$  contexts is followed by a unique set of  $k$  words. While this configuration is admittedly too symmetric to closely mimic natural language, it fulfills two functions. First, it allows for explicit calculation of the implicit geometry—specifically, the angles and norms of embeddings. Second, and more importantly for our analysis, it serves as a basis in Sec. 4.4.2 for approximating  $L^{\text{mm}}$  in more general scenarios, thereby motivating the heuristic proxy (P) in Sec. 4.4.3.

#### 4.4.1 Special case: support sets of equal sizes

**Proposition 2.** Fix  $k \in [V - 1]$  and suppose  $S$  contains all  $m = \binom{V}{k}$  support sets of size  $k$ . Then,

- (i) The logit matrix takes the form  $L^{\text{mm}} = (\mathbb{I}_V - \frac{1}{V}\mathbf{1}\mathbf{1}^\top)S \triangleq \tilde{S}$ .
- (ii) Word embeddings form equiangular tight frame (ETF) being equinorm and maximally separated.
- (iii) Context embeddings are equinorm and the embedding  $h_j$  of context  $j$  is co-linear to  $\sum_{z \in S_j} w_z$ .

In the symmetric setting of the proposition, we can analytically solve *NTP-SVM\** giving  $L^{\text{mm}}$  as expressed in the statement (i). In fact, it is possible to obtain an explicit characterization of the SVD of  $L^{\text{mm}}$ , which enables the precise definition of embeddings' geometries in statements (ii) and (iii); see also Prop. 5 in the appendix. For  $k = 1$  the embeddings recover the ETF geometry that has been previously shown for the setting of one-hot classification (Papayan et al., 2020). For general  $k > 1$ , word embeddings continue forming an ETF, but the geometry of context embeddings changes although they remain equinorm. Embedding  $h_j$  forms the same acute angle with all in-support word vectors  $w_v, v \in S_j$  and the same

obtuse angle with all out-of-support word vectors  $w_v, v \notin \mathcal{S}_j$ . Additionally, for  $k > 1$ , norms of word embeddings become larger than norms of context embeddings. See also Fig. 9 for a visualization of these properties.

#### 4.4.2 A simple candidate solution

Achieving an analytic expression for  $L^{\text{mm}}$ , as was possible in Prop. 2, might not always be possible. This challenge primarily arises from the combinatorial complexity of the constraints that adhere to the sparsity pattern of  $S$ . To mitigate this issue and potentially circumvent the need to solve the semidefinite program in  $\text{NTP-SVM}_*$ —which becomes computationally intensive for large  $V$  and  $m$ —we introduce a strategy that allows for a numerical verification of whether  $\tilde{S}$  of Prop. 2 solves  $\text{NTP-SVM}_*$ . Note that  $\tilde{S}$  readily satisfies the feasibility conditions for  $\text{NTP-SVM}_*$ , meeting all inequality constraints with equality. Below, we further introduce a dual-certificate condition that, when fulfilled by  $\tilde{S}$ , ensures optimality.

**Proposition 3.** *Let  $\tilde{S} = \mathbf{U}\Sigma\mathbf{V}^\top$  denote the SVD of  $\tilde{S} = (\mathbb{I}_V - \frac{1}{V}\mathbf{1}\mathbf{1}^\top)S$ . Define,  $A := \mathbf{U}\mathbf{V}^\top$ . If  $\forall j \in [m]$  and all  $v \notin \mathcal{S}_j$ , it holds that  $A[v, j] < 0$ , then  $\tilde{S}$  solves  $\text{NTP-SVM}_*$ .*

Verifying the certificate’s conditions only requires an SVD of  $\tilde{S}$ . In situations where these SVD factors can be analytically determined, the certificate also enables a formal proof of optimality (e.g., Prop. 2).

In more complex settings than that of Prop. 2, whether  $\tilde{S}$  satisfies the dual certificate in Prop. 3 or not, and thus its optimality, depends on the sparsity pattern of  $S$ . App. C.2 compares  $\tilde{S}$  to the optimal solution on a few small-scale examples. Although not necessarily optimal for all configurations of the support set, we conjecture that  $\tilde{S}$  gives an approximation with almost similar geometric properties of the optimal solution  $L^{\text{mm}}$ . This conjecture motivates a more efficient proxy for predicting the implicit geometry in the next section.

#### 4.4.3 A proxy for $\text{NTP-SVM}_*$ solutions

Considering  $\tilde{S} = (\mathbb{I}_V - \frac{1}{V}\mathbf{1}\mathbf{1}^\top)S$ , i.e., the column-wise centered support set, as a close proxy of the optimal solution of  $\text{NTP-SVM}_*$ , we provide a connection between the support sets and embeddings geometry. Suppose  $L^{\text{mm}} \approx \tilde{S}$ . Then, by Thm. 1,

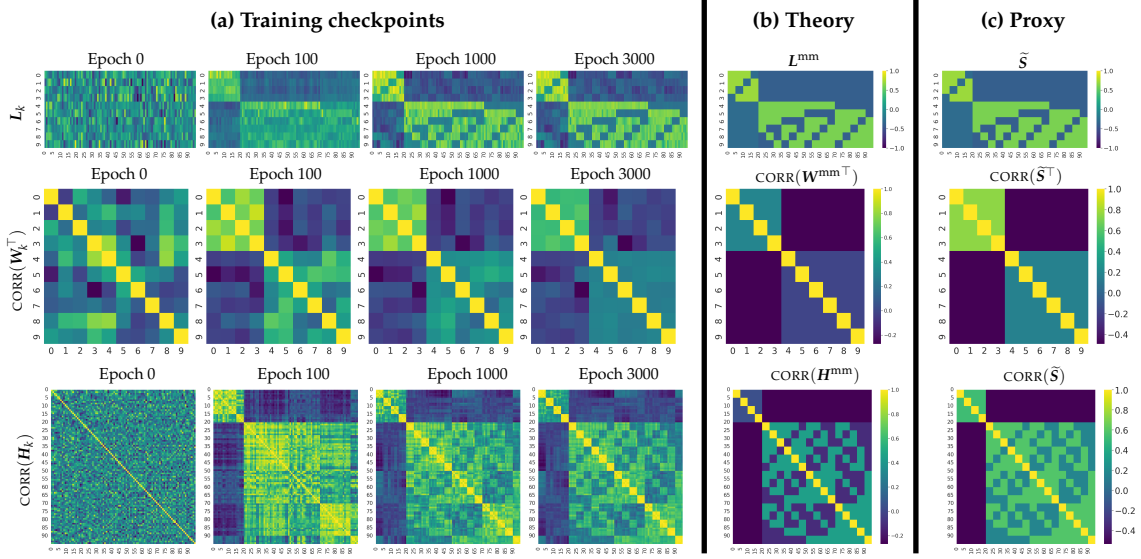
$$\begin{aligned} \mathbf{W}^{\text{mm}}\mathbf{W}^{\text{mm}\top} &= (\mathbf{L}^{\text{mm}}\mathbf{L}^{\text{mm}\top})^{\frac{1}{2}} \approx (\tilde{S}\tilde{S}^\top)^{\frac{1}{2}}, \\ \mathbf{H}^{\text{mm}\top}\mathbf{H}^{\text{mm}} &= (\mathbf{L}^{\text{mm}\top}\mathbf{L}^{\text{mm}})^{\frac{1}{2}} \approx (\tilde{S}^\top\tilde{S})^{\frac{1}{2}}. \end{aligned} \quad (11)$$

Computing the matrix square roots in Eq. (11) is as expensive as an SVD decomposition which can be prohibitive for large matrices. Instead, we suggest the following more computationally efficient proxies for estimating the directional component of the implicit geometries:


$$\begin{aligned} \mathbf{W}^{\text{mm}}\mathbf{W}^{\text{mm}\top} &\approx \tilde{S}\tilde{S}^\top = (\mathbb{I}_V - \frac{1}{V}\mathbf{1}_V\mathbf{1}_V^\top)\mathbf{S}\mathbf{S}^\top(\mathbb{I}_V - \frac{1}{V}\mathbf{1}_V\mathbf{1}_V^\top), \\ \mathbf{H}^{\text{mm}\top}\mathbf{H}^{\text{mm}} &\approx \tilde{S}^\top\tilde{S} = \mathbf{S}^\top(\mathbb{I}_V - \frac{1}{V}\mathbf{1}_V\mathbf{1}_V^\top)\mathbf{S}, \end{aligned} \quad (12)$$

This leads to Proxy (P) introduced in Sec. 1.2.

Ignoring for simplicity the projection to the  $\mathbb{I}_V - \frac{1}{V}\mathbf{1}_V\mathbf{1}_V^\top$  subspace, Eq. (12) gives a simple explanation for similarity between embeddings: For given context  $j \in [m]$ , the corresponding embedding has higher correlation with contexts  $j'$  whose support set intersection  $|\mathcal{S}_j \cap \mathcal{S}_{j'}|$  is larger. Similarly, word embeddings are closer if the respective words appear together in more support sets: word embedding  $w_z$  is more similar to  $w_{z'}$  when  $|\{j : w_z, w_{z'} \in \mathcal{S}_j, j \in [m]\}|$  is larger. The projection then fixes the bias of the parameters: since ridge-regularization is strongly convex and softmax is invariant to shift, any minimizer of  $\text{NTP-UFM}$  satisfies  $\mathbf{1}_V^\top\mathbf{W} = 0$  (see Lemma 2).



**Figure 4:** Evolution of **NTP-UFM** parameters  $L_k$ ,  $\text{CORR}(W_k^\top)$  and  $\text{CORR}(H_k)$  when training close to convergence to the empirical entropy  $\mathcal{H}$  (See Fig. 5). At the end of the training, the parameters align with the prediction of Thm. 1 ((a) vs (b)). Additionally, the correlation patterns between the embeddings closely follow the similarities between the support sets ((a) vs (c)). See Sec. 4.5 for details.



**Figure 5:** Numerical experiments on **NTP-UFM**. (a): CE converges closely to the empirical entropy  $\mathcal{H}$ , (b): Norms of the parameters grow, (c): While the parameters converge directionally to  $L^{\text{mm}}$  (Fig. 4), the projection of  $L_k$  on  $\mathcal{F}$  converges to the finite component  $L^{\text{in}}$  specified by the soft-labels  $\hat{p}_j$ , (d): Structural correlation between the learned embeddings and the proxy, i.e.,  $\text{SIM}(H, \tilde{S})$ ,  $\text{SIM}(W^\top, \tilde{S}^\top)$ . The correlation with  $H^{\text{mm}}$  and  $W^{\text{mm}\top}$  (instead of  $\tilde{S}$ ,  $\tilde{S}^\top$ ) is displayed with dashed lines for reference.

#### 4.5 Numerical simulation: Verifying the theory

Through numerical experiments training on purely synthetic data, we confirm that the solution derived from (stochastic) gradient descent optimization of **NTP-UFM** is consistent with the analysis presented in Sec. 4.

Specifically, we choose vocabulary size  $V = 10$ , embedding dimension  $d = 10$ , and  $m = 95$  training samples and generate (centered) support matrix  $\tilde{S}$  as shown in Fig. 4-(c). We also set  $\hat{\pi}_j = 1/m$  and for this fixed support matrix  $S$ , we generate random soft labels  $\hat{p}_j \in \Delta^{V-1}$  such that  $\hat{p}_{j,v} = 0, v \notin \mathcal{S}_j$ .

We display the training evolution in Fig. 5. We train **NTP-UFM** with SGD and small weight decay  $\lambda = 10^{-5}$  until reaching the empirical entropy lower-bound  $\mathcal{H} \approx 1.04$  within an order of  $10^{-4}$  as shown in Fig. 5-(a). This decrease in the loss is accompanied by a consistent increase in the parameter norms, as displayed in Panel (b). Despite the increase in norms, the projection of  $L_k$  onto the subspace  $\mathcal{F}$  remains close to the finite component  $L^{\text{in}}$  specified by Thm. 2. This is illustrated in Panel (c) and ensures the recovery of the soft-labels on the in-support tokens.

In Fig. 4, we visualize the logits  $L_k = W_k H_k$ , and cosine similarities  $\text{CORR}(W_k^\top)$  and  $\text{CORR}(H_k)$  of word and context embeddings, respectively. As shown in Panel (a) of Fig. 4, as training continues the parameters recover the implicit geometry in Panel (b), which visualizes the prediction made by Thm. 1 and Cor. 1. In Panel (c), we also visualize the cosine similarity of  $\tilde{S}$  and  $\tilde{S}^\top$ , which we introduced as a less expensive proxy of  $W^{\text{mm}}$  and  $H^{\text{mm}}$  (see Sec. 4.4.3). While the theory and proxy are not exactly the same, their heatmaps display similar structure. To verify the close relationship between the embeddings and the proxies quantitatively, we measure  $\text{SIM}(H, \tilde{S})$  and  $\text{SIM}(W^\top, \tilde{S}^\top)$ , where  $\text{SIM}(X, Y)$  measures the structural correlation between the two matrices (see Sec. 5.3). Fig. 5-(d) confirms the high correlation between the proxy (P) and the embeddings’ implicit geometry. We also note that in the special case where the support sets of two contexts are identical, the heatmaps also verify the subspace collapse Claim (C3): Along the diagonal blocks of  $\text{CORR}(H)$ , where the contexts with similar support sets lie, the samples have maximum correlation and are aligned.

We finally note that in this experiment  $L^{\text{mm}} = \tilde{S}$  (we verified that the dual-certificate condition in Prop. 3 holds).

## 5 Experiments

We now empirically validate our analysis on text data. We start with two small-scale synthetic datasets to examine Claims (C1)-(C4). We then experiment with a larger-scale dataset, where SVD calculation are computationally expensive, we examine the Proxy (P).

### 5.1 Datasets

For a detailed description of the datasets, refer to App. D.1.

**Synthetic.** We manually create simple (context, next-token) pairs, with context length of size  $T - 1 = 5$ . The dataset consists of  $n = 116$  samples, containing  $m = 16$  distinct contexts, and a vocabulary size of  $V = 30$ . Each context has a fixed support set length of  $S_j = 3$  and the empirical entropy lower-bound is  $\mathcal{H} = 1.6597$ .

**Simplified TinyStories.** For a more realistic but still controlled setup, we curate the dataset from the TinyStories corpus: We derive contexts  $\tilde{x}_j$  and support sets  $\mathcal{S}_j$  by choosing the most frequent word-level contexts with length  $T - 1 = 5$ . Here,  $V = 104$ ,  $n \approx 3050$  and  $m \approx 400$  with the empirical entropy lower-bound  $\mathcal{H} = 1.0166$ .

**TinyStories.** For more standard data, we use 100 stories sampled from TinyStories. Unlike the other two datasets, we do not sample over the contexts and support sets. We use a fixed context length  $T - 1 = 6$  for training. In the final dataset,  $V = 128$ ,  $m \sim 10^5$ , and the empirical entropy lower-bound is  $\mathcal{H} = 0.3112$ . We choose a small context window and vocabulary size to make tracking the distinct contexts and their support sets computationally manageable.

### 5.2 Models


We train decoder-only TF with 4 layers and  $d = 128$  for Synthetic and Simplified TinyStories, and 12 layers and  $d = 256$  for TinyStories. For training details, see App. D.1.

We verify two claims made in Sec. 4: 1) We assess whether numerically optimizing both the TF and NTP-UFM in each setup yields the same embedding geometries, 2) We evaluate the global solution of NTP-UFM, as specified by Thm. 1 and Cor. 1, and check whether it respects the same implicit geometry as the TF and the NTP-UFM. For the first one, we train NTP-UFM directly on the label distribution of each dataset; this is computationally feasible only for the first two datasets. For the second claim, we find the global solution  $L^{\text{mm}}$  of NTP-SVM, on each dataset using CVXPY (Diamond & Boyd, 2016) and  $H^{\text{mm}}$ ,  $W^{\text{mm}}$  using Cor. 1. If not computationally possible, we only use the Proxy (P) as reference.


### 5.3 Metrics.

We define the following metrics for verifying our theoretical results:





**Figure 6:** Experiments on Synthetic and Simplified TinyStories datasets. **(a):** CE approaches  $\mathcal{H}$ , **(b):** Parameters’ norms grow during training. **(c):**  $L_k$ ’s projection on the data subspace  $\mathcal{F}$  converges to the sparse component  $L^{\text{in}}$  specified by the soft-labels  $\hat{p}_j$ . **(d):**  $\text{SIM}(H_k, \tilde{S})$  (solid) and  $\text{SIM}(W_k^\top, \tilde{S}^\top)$  (dashed). At the final stage the learned embeddings exhibit high similarity with proxy **(P)**. **(e):** Same as (d), this time comparing with theory, i.e.,  $\text{SIM}(H_k, H^{\text{mm}})$  and  $\text{SIM}(W_k^\top, W^{\text{mm}\top})$ . Details in Sec. 5.



**Figure 7:** Similar to Fig. 5 and 6, this time on a 12-layer TF trained on a subset of 100 stories from the TinyStories dataset.

**Visualization of embeddings’ geometry.** To visualize the geometry of the word and context embeddings, we plot the heatmap of the normalized Gram matrix  $\text{CORR}(W^\top)$  and  $\text{CORR}(H)$ , where for matrix  $X$ ,  $\text{CORR}(X)$  is a matrix whose  $(i, j)$ -th entry is the cosine similarity between the respective columns of  $X$ , i.e.,

$$[\text{CORR}(X)]_{i,j} = \mathbf{x}_i^\top \mathbf{x}_j / (\|\mathbf{x}_i\| \|\mathbf{x}_j\|).$$

To check Proxy **(P)**, we compare the context and word embeddings with  $\text{CORR}(\tilde{S})$  and  $\text{CORR}(\tilde{S}^\top)$ .

**Quantifying geometric similarity.** To measure the structural similarity between  $W, H$  and  $\tilde{S}^\top, \tilde{S}$ , we use

$$\text{SIM}(X, Y) = (\sigma_{XY} + \epsilon) / (\sigma_X \sigma_Y + \epsilon),$$

as the correlation matrix between two matrices.<sup>5</sup> Here,  $\sigma_{XY}$  is the covariance, and  $\sigma_X$  and  $\sigma_Y$  are the standard deviations of  $X$  and  $Y$ , respectively, and  $\epsilon$  is a small constant for stable division. A value of  $\text{SIM}(X, Y) = 1$  indicates perfect structural correlation. To simplify the presentation, and with some abuse of notation, we denote the similarity metric for context and word embeddings as follows,

$$\text{SIM}(H, \tilde{S}) := \text{SIM}(\text{CORR}(H), \text{CORR}(\tilde{S})), \quad \text{SIM}(W^\top, \tilde{S}^\top) := \text{SIM}(\text{CORR}(W^\top), \text{CORR}(\tilde{S}^\top)).$$

**Recovery of soft-labels.** To verify Thm. 2, we measure the distance  $\|\mathcal{P}_{\mathcal{F}}(L_k) - L^{\text{in}}\|$  of  $L^{\text{in}}$  from the projection of  $L_k$  onto the subspace  $\mathcal{F}$ .

## 5.4 Results

In Sec. 1.1, we discussed Fig. 1, where we compared the logits, word and context embeddings trained on Simplified TinyStories by (a) TF, (b) NTP-UFM, and those predicted

<sup>5</sup>This is the Structural Similarity Index Measure (SSIM) used to measure the similarity between two images (Wang et al., 2004).

by (c) Claims **(C1)-(C3)**, and (d) Proxy **(P)**. Similar observations for the TinyStories and Synthetic datasets are shown in Figs. 3 and 13. For visualization details, see App. D.1.

Next, we discuss Fig. 6 that tracks our metrics during training for the Synthetic and Simplified TinyStories datasets. Fig. 6-(a) confirms that the loss is close to the empirical entropy  $\mathcal{H}$ . Fig. 6-(b) confirms the norm growth of the parameters. Note the slower growth of the context embeddings  $H$  in the TF model, which we suspect is due to the layer normalization at the final layer. We also confirm Claim **(C4)**, the recovery of the soft labels, in Fig. 6-(c). At last, Fig. 6-(d) and (e) track the correlation of the learned geometries with the proxy **(P)** and with the theoretical predictions  $(H^{mm}, W^{mm})$ . Note that all similarity metrics increase towards 1 as training progresses. In Figs. 14 and 15 in the appendix, we display the evolution of the embeddings by capturing their heatmaps at various training checkpoints. Interestingly, the final patterns emerge early in the training process. Fig. 7 displays analogous plots for TinyStories.

Finally, in Figs. 8 and 16, we visualize the text values of the contexts and words along with their geometry heatmaps. Observe that the embeddings’ similarities are consistent with the linguistic properties of words/contexts: Embeddings of contexts such as “boy named Timmy. Timmy” and “kid called Lilly. She” align closely as they are likely to be followed by similar set of words (i.e., they have similar sparsity patterns). Also, word embeddings tend to cluster by their grammatical categories, with verbs and adjectives often showing high similarity. For instance, the verbs go, eat, sing, wear, play, and the adjectives clever, fat, furry, lively, little have high similarity. This contrasts with the maximally separated word embeddings in the unrealistic, perfectly symmetric setting analyzed in Section 4.4.1, for which word embeddings were maximally separated. Here, the sparsity patterns of natural language induces an intricate geometry of word embeddings that accurately reflects linguistic structures such as grammatical categories.

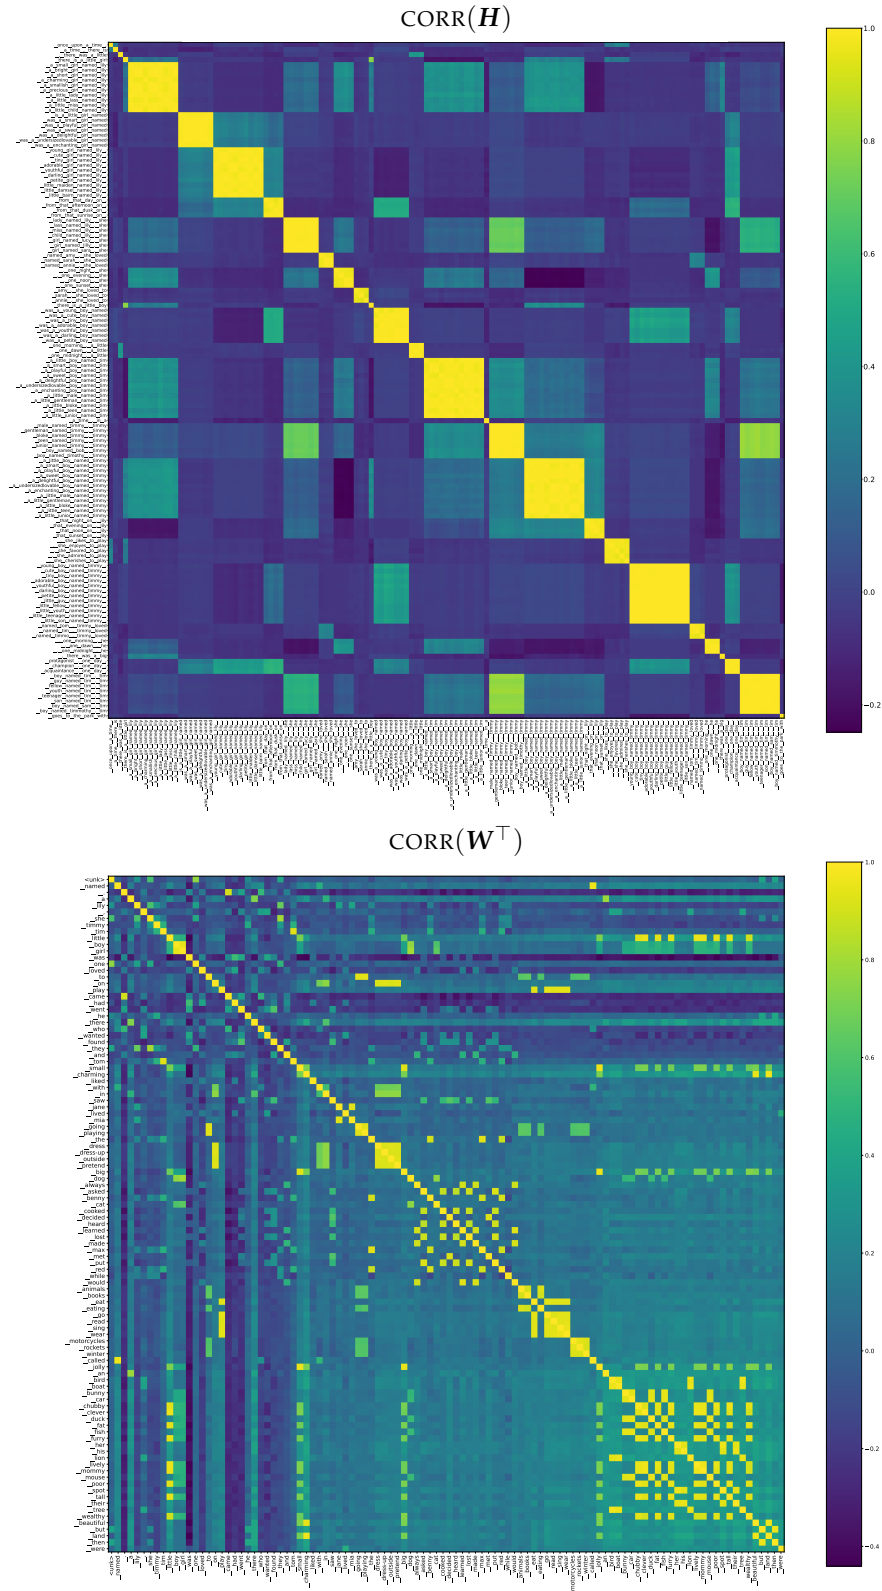
## 6 Concluding remarks: Limitations and future work

We have merged insights from the study of deep-net embeddings’ geometry in one-hot classification, implicit regularization theories, and a sparse soft-label classification framing of NTP, to arrive at a framework for examining word and context embeddings under the NTP paradigm. Key to our approach is the simplification to a problem of rank-constrained nuclear-norm-regularized CE minimization across logits (Eq. (8)). When the embedding dimension  $d$  is large relative to the vocabulary size  $V$ , we reveal how the similarities between word/context embeddings are heavily influenced by the sparsity patterns of the support sets. We do so by characterizing the implicit bias of NTP as it approaches the empirical entropy lower-bound leading to Claims **(C1)-(C4)**.

Looking forward, we see several promising avenues for expanding upon our framework and addressing its present limitations:

First, we believe it is feasible to broaden the regularization-path analysis to the non-convex realm of problem (8), potentially revealing a rank-constrained variant of **NTP-SVM\***. We hypothesize that the support matrices typical of natural language data might accommodate low-rank solutions that adhere to the constraints specified in Eqns. (9)-(10). This would allow us to relax the  $d \geq V$  constraint in our analysis; see App. D.4 for more discussions. But even under the condition  $d \geq V$  several exciting directions remain open. For example, investigating the implicit bias of GD—and potentially other optimizers like Adam, which are more commonly used in language model training but might present additional challenges—is important. In optimizing (8), does GD follow the regularization path? If not, to what extent does the solution learned by GD deviate from the sparse plus low-rank structure?

Second, while our initial focus was on establishing the framework, conducting preliminary analyses and proof-of-concept experiments, we encourage more extensive experimental exploration, especially with larger datasets and more complex model architectures with larger context window. Experiments in this paper focused on TF models due to their recent success in language. However, our theoretical results are independent of network design. Instead, they require the model to be over-parameterized enough to reach the entropy lower bound and ensure expressiveness in the embedding space. For an example, see App. D.2.



**Figure 8:** Implicit geometry of context (Top) and word (Bottom) embeddings and their associated text values for Simplified TinyStories dataset. Lighter color indicates higher similarity in the embedding space.

Third, it is intriguing to investigate how the distributional characteristics of word/context pairs in natural language, represented by the probability matrix  $P$ —and primarily by its support  $S$ —influence key properties like low-rankness and symmetry in the SVD factors of  $L^{\text{mm}}$ . Relatedly, we aim to leverage these insights to more effectively explore the empirical connection between  $L^{\text{mm}}$  and the centered support matrix  $\tilde{S}$  complementing the heuristic arguments in Sec. 4.4.3 in support of proxy (P). Considering this proxy, a low-rank structure of  $\tilde{S}$  would also imply a low-rank nature for  $L^{\text{mm}}$  (currently only indirectly supported by low-rank minimization). Fourth, although the framework encompasses full autoregressive training, our experimental study has focused on prediction of the last token prediction, with the exception of preliminary experiments in App. D.3. Deeper studies are required with sequences of variable length to empirically evaluate the brittleness of our model’s assumption, which enumerates distinct contexts regardless of their length or constituent token composition.

Fifth, our framework characterizes the geometry of word/context representations in a saturating regime where the NTP loss reaches its entropy lower-bound, suggesting that the model is sufficiently overparameterized. Although current state-of-the-art models do not reach this saturation regime and are likely not overparameterized, emerging evidence points to the need for extended training periods to achieve better generalization (Power et al., 2022) or to learn complex behaviors like hierarchical structures (Murty et al., 2023) and high-order in-context statistics (Edelman et al., 2024). It’s also plausible that as models grow larger and training methods become more efficient, extended training periods may become feasible. Investigating the potential benefits of prolonged training could yield valuable insights, and we hypothesize that the implicit optimization biases of NTP, as discussed in Thrampoulidis (2024) and our work, offer a compelling framework for such exploration. Additionally, determining the degree of overparameterization required to achieve the entropy lower-bound presents an intriguing challenge in its own right (Madden et al. (2024)).

Ultimately, our goal is to spur further exploration into the geometries of context and word embeddings. Akin to recent promising studies in image-classification that leverage the neural-collapse geometry, insights gained from such studies within NTP could inform the development of more effective loss functions, accelerate training processes, provide defenses against adversarial attacks, and guide the architectural choices for downstream tasks. More broadly, we aim to foster interest in deeper investigations into how the inductive bias of NTP influences the learning of linguistic patterns and regularities.

## Acknowledgments

This work was funded by the NSERC Discovery Grant No. 2021-03677, the Alliance Grant ALLRP 581098-22, and an Alliance Mission Grant. TB gratefully acknowledges the support from UBC’s Four Year Doctoral Fellowship and a scholarship from the Advanced Machine Learning Training Network at UBC. VV gratefully acknowledges support from BC Graduate Scholarship (BCGS) and Canadian Graduate Scholarship for Master’s (CGSM).

## References

- Carl Allen and Timothy Hospedales. Analogies explained: Towards understanding word embeddings. In *International Conference on Machine Learning*, pp. 223–231. PMLR, 2019.
- Sanjeev Arora, Yuanzhi Li, Yingyu Liang, Tengyu Ma, and Andrej Risteski. A latent variable model approach to pmi-based word embeddings. *Transactions of the Association for Computational Linguistics*, 4:385–399, 2016.
- Navid Azizan, Sahin Lale, and Babak Hassibi. Stochastic mirror descent on overparameterized nonlinear models. *IEEE Transactions on Neural Networks and Learning Systems*, 33(12): 7717–7727, 2021.

- Marco Baroni, Georgiana Dinu, and Germán Kruszewski. Don't count, predict! a systematic comparison of context-counting vs. context-predicting semantic vectors. In *Proceedings of the 52nd Annual Meeting of the Association for Computational Linguistics (Volume 1: Long Papers)*, pp. 238–247, 2014.
- Tina Behnia, Ganesh Ramachandra Kini, Vala Vakilian, and Christos Thrampoulidis. On the implicit geometry of cross-entropy parameterizations for label-imbalanced data. In *International Conference on Artificial Intelligence and Statistics*, pp. 10815–10838. PMLR, 2023.
- Samy Bengio and Yoshua Bengio. Taking on the curse of dimensionality in joint distributions using neural networks. *IEEE Transactions on Neural Networks*, 11(3):550–557, 2000.
- Yoshua Bengio, Réjean Ducharme, and Pascal Vincent. A neural probabilistic language model. *Advances in neural information processing systems*, 13, 2000.
- Daniel Biś, Maksim Podkorytov, and Xiuwen Liu. Too much in common: Shifting of embeddings in transformer language models and its implications. In *Proceedings of the 2021 conference of the North American chapter of the Association for Computational Linguistics: Human Language Technologies*, pp. 5117–5130, 2021.
- Tom B. Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel M. Ziegler, Jeffrey Wu, Clemens Winter, Christopher Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners, 2020.
- Matias D Cattaneo, Jason M Klusowski, and Boris Shigida. On the implicit bias of adam. *arXiv preprint arXiv:2309.00079*, 2023.
- Steven Diamond and Stephen Boyd. Cvxpy: A python-embedded modeling language for convex optimization. *Journal of Machine Learning Research*, 17(83):1–5, 2016.
- Benjamin L Edelman, Ezra Edelman, Surbhi Goel, Eran Malach, and Nikolaos Tsilivis. The evolution of statistical induction heads: In-context learning markov chains. *arXiv preprint arXiv:2402.11004*, 2024.
- Ronen Eldan and Yuanzhi Li. Tinstories: How small can language models be and still speak coherent english? *arXiv preprint arXiv:2305.07759*, 2023.
- Kawin Ethayarajh. How contextual are contextualized word representations? comparing the geometry of bert, elmo, and gpt-2 embeddings. *arXiv preprint arXiv:1909.00512*, 2019.
- Kawin Ethayarajh, David Duvenaud, and Graeme Hirst. Towards understanding linear word analogies. *arXiv preprint arXiv:1810.04882*, 2018.
- Cong Fang, Hangfeng He, Qi Long, and Weijie J Su. Exploring deep neural networks via layer-peeled model: Minority collapse in imbalanced training. *Proceedings of the National Academy of Sciences*, 118(43), 2021.
- Maryam Fazel. *Matrix rank minimization with applications*. PhD thesis, PhD thesis, Stanford University, 2002.
- Quinn Fisher, Haoming Meng, and Vardan Papyan. Pushing boundaries: Mixup's influence on neural collapse. *arXiv preprint arXiv:2402.06171*, 2024.
- Jun Gao, Di He, Xu Tan, Tao Qin, Liwei Wang, and Tie-Yan Liu. Representation degeneration problem in training natural language generation models, 2019.
- Nathan Godey, Éric de la Clergerie, and Benoît Sagot. Anisotropy is inherent to self-attention in transformers. *arXiv preprint arXiv:2401.12143*, 2024.



- Michael Grant and Stephen Boyd. CVX: Matlab software for disciplined convex programming, version 2.1. <http://cvxr.com/cvx>, March 2014.
- Suriya Gunasekar, Jason Lee, Daniel Soudry, and Nathan Srebro. Characterizing implicit bias in terms of optimization geometry. In *International Conference on Machine Learning*, pp. 1832–1841. PMLR, 2018a.
- Suriya Gunasekar, Jason D Lee, Daniel Soudry, and Nati Srebro. Implicit bias of gradient descent on linear convolutional networks. *Advances in Neural Information Processing Systems*, 31:9461–9471, 2018b.
- Zellig S Harris. Distributional structure. *Word*, 10(2-3):146–162, 1954.
- Tatsunori B Hashimoto, David Alvarez-Melis, and Tommi S Jaakkola. Word embeddings as metric recovery in semantic spaces. *Transactions of the Association for Computational Linguistics*, 4:273–286, 2016.
- Brian Hutchinson, Mari Ostendorf, and Maryam Fazel. A sparse plus low rank maximum entropy language model. In *INTERSPEECH*, pp. 1676–1679, 2012.
- Brian Hutchinson, Mari Ostendorf, and Maryam Fazel. Exceptions in language as learned by the multi-factor sparse plus low-rank language model. In *2013 IEEE International Conference on Acoustics, Speech and Signal Processing*, pp. 8580–8584. IEEE, 2013.
- Brian Hutchinson, Mari Ostendorf, and Maryam Fazel. A sparse plus low-rank exponential language model for limited resource scenarios. *IEEE/ACM Transactions on Audio, Speech, and Language Processing*, 23(3):494–504, 2015.
- Ziwei Ji and Matus Telgarsky. Risk and parameter convergence of logistic regression. *arXiv preprint arXiv:1803.07300*, 2018.
- Ziwei Ji and Matus Telgarsky. Directional convergence and alignment in deep learning. *Advances in Neural Information Processing Systems*, 33:17176–17186, 2020.
- Ziwei Ji, Miroslav Dudík, Robert E Schapire, and Matus Telgarsky. Gradient descent follows the regularization path for general losses. In *Conference on Learning Theory*, pp. 2109–2136. PMLR, 2020.
- Daniel Jurafsky and James H. Martin. *Speech and Language Processing*. Prentice Hall, 3 edition, 2023. Draft.
- Omer Levy and Yoav Goldberg. Neural word embedding as implicit matrix factorization. *Advances in neural information processing systems*, 27, 2014.
- Bohan Li, Hao Zhou, Junxian He, Mingxuan Wang, Yiming Yang, and Lei Li. On the sentence embeddings from pre-trained language models. *arXiv preprint arXiv:2011.05864*, 2020.
- Pengyu Li, Yutong Wang, Xiao Li, and Qing Qu. Neural collapse in multi-label learning with pick-all-label loss. *arXiv preprint arXiv:2310.15903*, 2023.
- Hong Liu, Sang Michael Xie, Zhiyuan Li, and Tengyu Ma. Same pre-training loss, better downstream: Implicit bias matters for language models. In *International Conference on Machine Learning*, pp. 22188–22214. PMLR, 2023.
- Kaifeng Lyu and Jian Li. Gradient descent maximizes the margin of homogeneous neural networks. In *International Conference on Learning Representations*, 2020.
- Liam Madden, Curtis Fox, and Christos Thrampoulidis. Upper and lower memory capacity bounds of transformers for next-token prediction. *arXiv preprint arXiv:2405.13718*, 2024.
- Tomas Mikolov, Kai Chen, Greg Corrado, and Jeffrey Dean. Efficient estimation of word representations in vector space. *arXiv preprint arXiv:1301.3781*, 2013a.

- Tomas Mikolov, Ilya Sutskever, Kai Chen, Greg S Corrado, and Jeff Dean. Distributed representations of words and phrases and their compositionality. *Advances in neural information processing systems*, 26, 2013b.
- Tomáš Mikolov, Wen-tau Yih, and Geoffrey Zweig. Linguistic regularities in continuous space word representations. In *Proceedings of the 2013 conference of the north american chapter of the association for computational linguistics: Human language technologies*, pp. 746–751, 2013c.
- Dustin G Mixon, Hans Parshall, and Jianzong Pi. Neural collapse with unconstrained features. *arXiv preprint arXiv:2011.11619*, 2020.
- Jiaqi Mu, Suma Bhat, and Pramod Viswanath. All-but-the-top: Simple and effective post-processing for word representations. *arXiv preprint arXiv:1702.01417*, 2017.
- Shikhar Murty, Pratyusha Sharma, Jacob Andreas, and Christopher D Manning. Grokking of hierarchical structure in vanilla transformers. *arXiv preprint arXiv:2305.18741*, 2023.
- Mor Shpigel Nacson, Jason Lee, Suriya Gunasekar, Pedro Henrique Pamlona Savarese, Nathan Srebro, and Daniel Soudry. Convergence of gradient descent on separable data. In *The 22nd International Conference on Artificial Intelligence and Statistics*, pp. 3420–3428. PMLR, 2019.
- Vardan Papyan, XY Han, and David L Donoho. Prevalence of neural collapse during the terminal phase of deep learning training. *Proceedings of the National Academy of Sciences*, 117(40):24652–24663, 2020.
- Jeffrey Pennington, Richard Socher, and Christopher D Manning. Glove: Global vectors for word representation. In *Proceedings of the 2014 conference on empirical methods in natural language processing (EMNLP)*, pp. 1532–1543, 2014.
- Scott Pesme, Loucas Pillaud-Vivien, and Nicolas Flammarion. Implicit bias of sgd for diagonal linear networks: a provable benefit of stochasticity. *Advances in Neural Information Processing Systems*, 34:29218–29230, 2021.
- Alethea Power, Yuri Burda, Harri Edwards, Igor Babuschkin, and Vedant Misra. Grokking: Generalization beyond overfitting on small algorithmic datasets. *arXiv preprint arXiv:2201.02177*, 2022.
- Alec Radford, Karthik Narasimhan, Tim Salimans, Ilya Sutskever, et al. Improving language understanding by generative pre-training. 2018.
- Alec Radford, Jeffrey Wu, Rewon Child, David Luan, Dario Amodei, Ilya Sutskever, et al. Language models are unsupervised multitask learners. *OpenAI blog*, 1(8):9, 2019.
- Benjamin Recht, Maryam Fazel, and Pablo A Parrilo. Guaranteed minimum-rank solutions of linear matrix equations via nuclear norm minimization. *SIAM review*, 52(3):471–501, 2010.
- Saharon Rosset, Ji Zhu, and Trevor Hastie. Margin maximizing loss functions. In *NIPS*, pp. 1237–1244, 2003.
- Nikunj Saunshi, Sadhika Malladi, and Sanjeev Arora. A mathematical exploration of why language models help solve downstream tasks. *arXiv preprint arXiv:2010.03648*, 2020.
- Hinrich Schütze, Christopher D Manning, and Prabhakar Raghavan. *Introduction to information retrieval*, volume 39. Cambridge University Press Cambridge, 2008.
- Claude Elwood Shannon. A mathematical theory of communication. *The Bell system technical journal*, 27(3):379–423, 1948.
- Daniel Soudry, Elad Hoffer, Mor Shpigel Nacson, Suriya Gunasekar, and Nathan Srebro. The implicit bias of gradient descent on separable data. *The Journal of Machine Learning Research*, 19(1):2822–2878, 2018.

- Nathan Srebro, Jason Rennie, and Tommi Jaakkola. Maximum-margin matrix factorization. *Advances in neural information processing systems*, 17, 2004.
- Haoyuan Sun, Kwangjun Ahn, Christos Thrampoulidis, and Navid Azizan. Mirror descent maximizes generalized margin and can be implemented efficiently. *Advances in Neural Information Processing Systems*, 35:31089–31101, 2022.
- Davoud Ataee Tarzanagh, Yingcong Li, Christos Thrampoulidis, and Samet Oymak. Transformers as support vector machines, 2023a.
- Davoud Ataee Tarzanagh, Yingcong Li, Xuechen Zhang, and Samet Oymak. Max-margin token selection in attention mechanism, 2023b.
- Christos Thrampoulidis. Implicit bias of next-token prediction. *arXiv preprint arXiv:2402.18551*, 2024.
- Christos Thrampoulidis, Ganesh R Kini, Vala Vakilian, and Tina Behnia. Imbalance trouble: Revisiting neural-collapse geometry. *arXiv preprint arXiv:2208.05512*, 2022.
- William Timkey and Marten Van Schijndel. All bark and no bite: Rogue dimensions in transformer language models obscure representational quality. *arXiv preprint arXiv:2109.04404*, 2021.
- Joseph Turian, Lev Ratinov, and Yoshua Bengio. Word representations: a simple and general method for semi-supervised learning. In *Proceedings of the 48th annual meeting of the association for computational linguistics*, pp. 384–394, 2010.
- Gal Vardi and Ohad Shamir. Implicit regularization in relu networks with the square loss. In *Conference on Learning Theory*, pp. 4224–4258. PMLR, 2021.
- Zhou Wang, Alan C Bovik, Hamid R Sheikh, and Eero P Simoncelli. Image quality assessment: from error visibility to structural similarity. *IEEE transactions on image processing*, 13(4):600–612, 2004.
- Robert Wu and Vardan Papyan. Linguistic collapse: Neural collapse in (large) language models. *arXiv preprint arXiv:2405.17767*, 2024.
- Zhilin Yang, Zihang Dai, Ruslan Salakhutdinov, and William W Cohen. Breaking the softmax bottleneck: A high-rank rnn language model. *arXiv preprint arXiv:1711.03953*, 2017.
- Zhihui Zhu, Tianyu Ding, Jinxin Zhou, Xiao Li, Chong You, Jeremias Sulam, and Qing Qu. A geometric analysis of neural collapse with unconstrained features. *Advances in Neural Information Processing Systems*, 34, 2021.

# Contents

<b>1</b>	<b>Introduction</b>	<b>1</b>
1.1	Methodology	3
1.2	Summary of findings	4
<b>2</b>	<b>Related work</b>	<b>6</b>
<b>3</b>	<b>Formulation</b>	<b>6</b>
3.1	NTP objective as sparse soft-label classification	7
3.2	Unconstrained features model for NTP training	7
3.3	Reformulation in terms of logits	8
<b>4</b>	<b>Analysis of unconstrained features model for NTP training</b>	<b>8</b>
4.1	NTP-SVM logits	8
4.2	Regularization-path analysis of logits	9
4.3	Geometry of embeddings	10
4.4	On the solution of NTP-SVM <sub>*</sub>	10
4.4.1	Special case: support sets of equal sizes	10
4.4.2	A simple candidate solution	11
4.4.3	A proxy for NTP-SVM <sub>*</sub> solutions	11
4.5	Numerical simulation: Verifying the theory	12
<b>5</b>	<b>Experiments</b>	<b>13</b>
5.1	Datasets	13
5.2	Models	13
5.3	Metrics	13
5.4	Results	14
<b>6</b>	<b>Concluding remarks: Limitations and future work</b>	<b>15</b>
<b>A</b>	<b>Related work: Detailed version</b>	<b>24</b>
A.1	Word embeddings as matrix factorization	24
A.2	Overparameterization and Implicit bias/regularization of GD	25
A.3	Neural collapse geometry	26
A.4	Connections to Word2Vec	27
<b>B</b>	<b>Proofs</b>	<b>28</b>
B.1	Centering of word embeddings	28
B.2	Proof of Lemma 1	28
B.3	Proof of Theorems 1 and 2	30
B.3.1	Notations recap	30
B.3.2	When does the NTP loss attain its entropy lower-bound?	31
B.3.3	Ball-constrained CE minimization and main result	32
B.3.4	The in-support NTP loss and its minimizer $L^{\text{in}}$	33
B.3.5	Auxiliary lemmas	34
B.3.6	Proof of Theorem 3	35
B.4	Proof of Proposition 1	37
<b>C</b>	<b>Additional discussions on the solution of NTP-SVM<sub>*</sub></b>	<b>38</b>
C.1	Special case: support sets of equal sizes	38
C.2	A simple candidate solution	38
C.3	Proof of Proposition 3	39
C.4	Proof of Propositions 2 and 5	40
<b>D</b>	<b>Numerical experiments</b>	<b>42</b>
D.1	Additional details and results: Sec. 5	42
D.1.1	Datasets	42
D.1.2	Training details	43
D.1.3	Additional results	43

D.2 Other architectures . . . . .	47
D.3 Auto-regressive training . . . . .	47
D.4 Discussion on large vocabulary setting . . . . .	47



## A Related work: Detailed version

This section expands upon the section on related work in the main body.

### A.1 Word embeddings as matrix factorization

Neural probabilistic language models are designed to predict the probability distribution of a target word  $v \in \mathcal{V}$  within a given vocabulary  $\mathcal{V}$ , based on a sequence of preceding or surrounding words of (variable or fixed) length  $t$  known as the context  $z \in \mathcal{V}^t$ . In these models, word and context representations (aka embeddings) are parameterized by their internal architecture (Bengio et al., 2000). The emergence of energy-based models like the Word2Vec family, particularly through the introduction of training methodologies such as the Skip-Gram with Negative Sampling (SGNS) objective by Mikolov et al. (2013b;a), marked a significant advancement in learning high quality embeddings; e.g., see Jurafsky & Martin (2023). SGNS aims to optimize the alignment between vectors of *target* words and the *context* words surrounding the target, while disassociating those of randomly drawn pairs. This is achieved by optimizing a log-bilinear model that learns unconstrained embeddings for words in two roles: as targets and as contexts. While intuitive, the objective remained somewhat mysterious until the seminal analysis by Levy & Goldberg (2014) offered a formal examination of embeddings generated by the SGNS objective, framing it as a form of weighted matrix factorization. Specifically for embeddings of large dimensionality, Levy and Goldberg demonstrated that SGNS implicitly factorizes a shifted version of the Pointwise Mutual Information (PMI) matrix, denoted as **PMI**, which is a measure of association between words and contexts by having entries  $\text{PMI}[z, j] = \log \frac{\Pr(z|j)}{\Pr(z)} = \log \frac{\hat{p}_{j,z}}{\Pr(z)}$  for target word  $z$  and context word  $j$ . This result establishes a fundamental connection between count-based and predictive models (Pennington et al., 2014). However, this approach presumes non-zero occurrences of all context-word pairs in the training set, a condition not always met, leading to undefined factorization for pairs with zero co-occurrence (resulting in  $\text{PMI}[z, j] = -\infty$ ). To address the possibility of negative infinity values, an approximation involving a positive, sparse, and non-negative PMI (PPMI) matrix was proposed. While this adaptation offered an intriguing alternative to the original training objectives, the geometric characteristics of the resultant embeddings in comparison to those derived from the original loss remain unclear.

Our approach mirrors that of Levy & Goldberg (2014) to some extent (see also Pennington et al. (2014)), especially in treating NTP as a soft-label classification across distinct contexts and presuming unconstrained context embeddings (although these concepts are not clearly mentioned in those original works). At a very basic level our work distinguishes itself by focusing on the NTP objective, which, although akin to SGNS and NC objectives, employs cross-entropy loss instead of sigmoid functions. Aside from this distinction, we identify the following conceptual divergences, contributing to our novel perspective on this line of inquiry: **(i)** Firstly, we aim to characterize the geometry of embeddings as learned directly from the loss function, rather than through approximations. **(ii)** Secondly, although **NTP-UFM**, as a log-bilinear model, is in spirit to the log-bilinear models trained by Word2Vec, we only use such log-bilinear models as an analytically-tractable proxy to expressive deep neural probabilistic language models such as transformers. **(iii)** Thirdly, we confront the sparsity in the probabilistic labels of each context head-on, unlike other matrix-factorization methods that circumvent sparsity in a heuristic manner through smoothing or weighting techniques (Levy & Goldberg, 2014; Pennington et al., 2014). This helps us identify that sparsity leads to diverging embeddings, a scenario where setting the loss gradient to zero is infeasible with finite weights. **(iv)** Fourthly, by examining embeddings through the lens of the regularization path—a surrogate for gradient descent optimization—we unveil that word and context embeddings emerge from the factorization of a matrix  $L^{\text{in}} + RL^{\text{mm}}$ . Here,  $L^{\text{in}}$  captures the frequency of word-context occurrences akin to the PMI, setting non-occurring pairs to zero while retaining positive values for occurring pairs, much like the PPMI matrix. The directional component  $L^{\text{mm}}$ , becoming dominant as weights diverge, reflects the explicit sparsity patterns of the context-word probability matrix.

The above-mentioned conceptual novelties can also be directly applied to the SGNS objective generalizing the result of [Levy & Goldberg \(2014\)](#) to sparse language patterns. We show this in App. [A.4](#).

The seminal contribution of [Levy & Goldberg \(2014\)](#) has inspired numerous subsequent studies, leveraging the geometric insights of embeddings to uncover linguistic phenomena such as the linear relationships underlying word analogies ([Allen & Hospedales, 2019](#); [Hashimoto et al., 2016](#); [Arora et al., 2016](#); [Ethayarajh et al., 2018](#); [Mikolov et al., 2013c](#)). Other works, investigate the *anisotropy* of the learned word/context embeddings for both static models such as Word2Vec ([Mu et al., 2017](#)) and contextualized models such as pre-trained BERT and GPT-2 ([Ethayarajh, 2019](#); [Li et al., 2020](#)). This property of the embeddings has been attributed to the imbalanced long-tail distribution of language ([Gao et al., 2019](#); [Li et al., 2020](#)), shifts in the embedding space ([Biś et al., 2021](#)), and characteristic of the network architecture ([Godey et al., 2024](#)). Various techniques have been proposed to improve the isotropy of the embeddings, which yields embeddings with better language regularities (e.g., ([Mu et al., 2017](#); [Gao et al., 2019](#); [Biś et al., 2021](#); [Timkey & Van Schijndel, 2021](#))). We envision that our fresh perspective on geometric analysis and our findings could similarly motivate further research in this domain, extending the understanding and application of language model embeddings.

Finally, we note that a *sparse plus low-rank* structure for parameters learned with weighted exponential language models trained on fixed n-gram word/context features has been observed in [Hutchinson et al. \(2012\)](#). Through experiments, the authors decompose the weights learned by optimizing a weighted maximum entropy objective into sparse and low-rank components. They demonstrate that explicitly incorporating sparsity and low-rank constraints into their weighted maximum entropy framework can enhance both performance and interpretability ([Hutchinson et al., 2012; 2013](#)), particularly when training data is limited ([Hutchinson et al., 2015](#)). However, these components serve different roles compared to our work. In their framework, the low-rank component captures frequent word co-occurrences, while the sparse component records infrequent words and special lexical items. In contrast, our analysis shows via the regularization path framework that the NTP objective produces a low-rank component from a nuclear-norm max-margin problem between in-support and out-support tokens, with the sparse component encoding the frequency of the in-support tokens.

## A.2 Overparameterization and Implicit bias/regularization of GD

Recent studies on gradient-based optimization methods, such as SGD and its variants, have focused on overparameterized settings in empirical risk minimization without explicit regularization. This interest stems from observations that deep neural networks, especially in image classification, perform exceptionally well even with minimal or no weight decay. This research area, known as “implicit bias” or “implicit regularization” of GD emerged to explore which among the numerous potential minimizers in overparameterized settings are preferred by algorithms like GD. Notably, [Soudry et al. \(2018\)](#); [Ji & Telgarsky \(2018\)](#) showed that for linear logistic regression with one-hot labels, GD aligns with the hard-margin SVM solution, essentially becoming the max-margin classifier. Additionally, [Ji & Telgarsky \(2020\)](#) indicated that GD’s optimization path in linear models is analogous to the regularization path, which looks at risk minimizers as regularization diminishes. This path is analytically more accessible and offers a useful framework to study GD’s behavior. In our work, we apply this idea to investigate the regularization path of NTP training. These foundational results on implicit bias have spurred further research into stochastic and adaptive gradient methods ([Nacson et al., 2019](#); [Pesme et al., 2021](#); [Sun et al., 2022](#); [Azizan et al., 2021](#); [Cattaneo et al., 2023](#)), and more intricate architectures ([Lyu & Li, 2020](#); [Ji & Telgarsky, 2020](#); [Gunasekar et al., 2018a;b](#))) including recent studies on transformers ([Tarzanagh et al., 2023b;a](#)); for a review, refer to [Vardi & Shamir \(2021\)](#). These insights have also paved the way for generalization analysis, linking GD’s generalization to that of SVM-like solutions. This connection has been explored in high-dimensional contexts where traditional margin-based bounds are inadequate, revealing scenarios where no regularization yields optimal outcomes, a phenomenon known as ‘harmless interpolation’ or ‘benign overfitting.’ Most of these findings, whether on optimization or generalization, are confined to the one-hot classification framework.

In the language modeling literature, several theoretical works suggest that achieving the empirical entropy lower-bound at the pre-training stage results in better down-stream performance (Liu et al., 2023; Saunshi et al., 2020). Liu et al. (2023) attributes this better down-stream performance to the implicit bias of the pre-training algorithms, by showing that the optimizers prefer the more transferable models among the ones that achieve the same minimal loss value. A formal deviation from one-hot to soft-label classification for NTP is the work by (Thrampoulidis, 2024), which partially inspired our research. While their analysis, assuming fixed context embeddings, applies to linear models, our work extends this by optimizing both word and context embeddings, leading to a nonlinear model. We discover that this nonlinearity yields a preference towards minimizing the nuclear norm of the logits over the Euclidean norm of the word embedding matrix, as seen in (Thrampoulidis, 2024). Our theoretical contributions thus expand on their findings into the nonconvex domain.

### A.3 Neural collapse geometry

A third area of research closely related to our work is the investigation of neural-collapse (NC) geometry, which seeks to understand the geometry of last-layer features and classifier weights in deep networks optimized in the interpolating regime. Initiated by Papayan et al. (2020), this line of research uncovered a universal phenomenon across various architectures and datasets, where classifiers and last-layer features consistently converge to an equiangular tight frame (ETF). This convergence subsumes the aggregation of same-class features around their class mean, a phenomenon termed NC. The theoretical underpinnings for these empirical findings were later provided by analyzing the unconstrained features model (UFM) (Mixon et al., 2020; Fang et al., 2021; Zhu et al., 2021), suggesting that highly overparameterized networks can produce arbitrary last-layer features. This concept, reminiscent of the early framework used in Word2Vec models (Levy & Goldberg, 2014) as discussed above (see also Yang et al. (2017)), has since spurred extensive research in diverse contexts.

Among the vast body of work on NC geometry—an exhaustive reference of which is beyond our scope—, two studies stand out for their relevance to our research. Thrampoulidis et al. (2022) first adapted the UFM to describe geometries diverging from symmetric ETF by moving away from the balanced data assumption, demonstrating that logits optimizing the UFM with vanishing regularization align with a nuclear-norm SVM problem; see also Behnia et al. (2023). This approach also bridges NC geometry with the implicit bias literature previously discussed. Our work extends their principles to the NTP objective, differing from the one-hot encoding classification focus of their study. We take a similar approach, but this time applied to the NTP objective. This is unlike their result which only applies to one-hot encoding classification. In fact, our formulation in  $\text{NTP-SVM}_*$  extends their findings, encompassing them as a specific instance when each context is followed by only one token—a scenario less applicable to the linguistic contexts we focus on. Moreover, we demonstrate (see Theorem 2) that the NTP framework introduces a finite projection component to the regularization path, a feature absent in one-hot settings.

Our findings also resonate to certain extent with the recent multilabel NC geometry study in Li et al. (2023), where each example, typically an image, is tagged with multiple correct labels, represented as a  $k$ -hot vector. In contrast, our NTP framework, modeled as soft-label classification, associates each context with a probability vector. This distinction makes our framework more versatile, applicable beyond the restrictive case where soft-label classification aligns with multi-label classification—specifically, when all relevant tokens are equally likely. There are also crucial differences in the outcomes. Their analysis assumes an equal number of training samples of multiplicity  $k = 1$  for each category, a condition that, in our context, would necessitate every vocabulary word to be followed by an identical number of contexts that are only followed by that word—a significantly limiting constraint. Our findings, however, are not bound by this assumption.

In an attempt to extend this line of work to language setups, contemporaneous work Wu & Pappan (2024), through experiments, suggests a correlation between the degree of convergence to the NC geometry and the validation loss for different language models trained on TinyStories for a few epochs. Different to us, they implicitly treat the NTP

objective as one-hot classification and suggest that better *per-class* (i.e., next-token) collapse of the context embeddings leads to improved validation loss. In contrast, our analysis reveals a more intricate finding: within a sparse soft-label formulation of NTP training with language, it is not *per-class* collapse that minimizes the training objective. Instead, the objective is minimized when context embeddings with the same *support set* collapse after being projected onto *subspace*  $\mathcal{F}_\perp$ , which is defined by the sparsity pattern of the context-word co-occurrence matrix (see Sec. 4.3). Our approach also comes with an analysis framework rather than relying solely on experiments.

#### A.4 Connections to Word2Vec

We apply our geometry analysis to Word2Vec. Consider the following Skip-gram objective:

$$\frac{1}{T} \sum_{t \in [T]} \sum_{-c \leq j \leq c, j \neq 0} \log p(z_{t+j} | z_t) = \frac{1}{T} \sum_{t \in [T]} \sum_{-c \leq j \leq c, j \neq 0} \log p(\mathbf{S}_{z_{t+j}}(\mathbf{H} \mathbf{w}_{z_t})),$$

where  $\mathbf{H} \in \mathbb{R}^{V \times d}$  denotes the embedding matrix of *context* words  $z_{t+j}, j \neq 0$  and  $\mathbf{W} \in \mathbb{R}^{V \times d}$  denotes the embedding matrix of *target* words  $z_t$ . As presented in Mikolov et al. (2013c), skip-gram actually implements a heuristic of this objective that aims to simplify the computation of the normalizing factor in the softmax. Second, consider the CBOW objective:

$$\frac{1}{T} \sum_{t \in [T]} \log p(z_t | \{z_{t+j}\}_{\{|j| \in [c], j \neq 0\}}) = \frac{1}{T} \sum_{t \in [T]} \log p \left( \mathbf{S}_{z_t} \left( \mathbf{W} \cdot \frac{1}{2c-1} \sum_{-c \leq j \leq c, j \neq 0} \mathbf{h}_{z_{t+j}} \right) \right),$$

where  $\frac{1}{2c-1} \sum_{-c \leq j \leq c, j \neq 0} \mathbf{h}_{z_{t+j}}$  is the continuous-bags-of-words embedding of the context  $\{z_{t+j}\}_{\{|j| \in [c], j \neq 0\}}$ .

Suppose the CBOW model with context consisting of only the previous word. To keep notation consistent with the rest of the paper, let  $j \in [m] = [V]$  denote index of *distinct* contexts (aka previous words) and  $z \in [V]$  denote index of (next) words. Note the embedding matrix  $\mathbf{H} \in \mathbb{R}^{d \times V}$  is now of same dimension as  $\mathbf{W}^\top$ . Let  $\hat{\pi}_j$  denote the marginal of context words and  $\hat{p}_{j,z}$  denote the conditional probability of word  $z$  following context word  $j$ . Then, the CBOW objective becomes equivalent to

$$\sum_{j \in [V]} \hat{\pi}_j \sum_{z \in [V]} \hat{p}_{j,z} \log (\mathcal{S}_z(\mathbf{W} \mathbf{h}_j)).$$

This now fits exactly our framework and the results for NTP apply mutatis mutandis.

In Mikolov et al. (2013c), instead of minimizing the CE objective, they actually use the Negative-Sampling (SGNS) objective. Levy & Goldberg (2014) rewrites the SGNS objective as follows:

$$\sum_{j \in [m]} \sum_{z \in [V]} -\Pr(j, z) \log (\sigma(\mathbf{w}_z^\top \mathbf{h}_j)) - k \Pr(j) \Pr(z) \log (1 - \sigma(\mathbf{w}_z^\top \mathbf{h}_j)),$$

where  $\Pr(j, z) := \hat{\pi}_j \hat{p}_{j,z}$ ,  $\Pr(j) := \hat{\pi}_j$ ,  $\Pr(z) = \sum_{j \in [m]} \hat{\pi}_j \hat{p}_{j,z}$  and  $k \in \mathbb{N}$  is a hyperparameter. Defining

$$q_{j,z} := \frac{\Pr(j,z)}{\Pr(j) \Pr(z)} \in [0, 1],$$

we can rewrite the SGNS objective as

$$\sum_{j \in [m]} \sum_{z \in [V]} \Pr(j) \Pr(z) \left( k + \frac{\Pr(j,z)}{\Pr(j) \Pr(z)} \right) \left( -q_{j,z} \log (\sigma(\mathbf{w}_z^\top \mathbf{h}_j)) - (1 - q_{j,z}) \log (1 - \sigma(\mathbf{w}_z^\top \mathbf{h}_j)) \right).$$

For each  $(j, z)$ -pair, the rightmost term in the expression above is lower bounded by  $\mathcal{H}(q_{j,z}) := -q_{j,z} \log(q_{j,z}) - (1 - q_{j,z}) \log(1 - q_{j,z})$ .

By following the same arguments developed for NTP, it can be checked that the lower bound is attained if and only if the following conditions hold:

A. There exists logit matrix  $L_1 \in \mathbb{R}^{V \times m}$  (of rank  $\leq d$ ) such that

$$L_1[z, j] = \text{PMI}(z, j) - \log(k) = \log\left(\frac{\Pr(j, z)}{\Pr(j)\Pr(z)}\right) - \log(k), \quad \forall j \in [m], z \in \mathcal{S}_j.$$

B. There exists logit matrix  $L_2 \in \mathbb{R}^{V \times m}$  (of rank  $\leq d$ ) such that

$$\begin{aligned} L_2[z, j] &= 0, \quad \forall j \in [m], z \in \mathcal{S}_j, \\ L_2[v, j] &< 0, \quad \forall j \in [m], v \notin \mathcal{S}_j. \end{aligned}$$

The logit matrix  $L_1$  is the shifted PMI matrix as also characterized in [Levy & Goldberg \(2014, Eq. \(7\)\)](#). This component plays a similar role to the finite component  $L^{\text{in}}$  in the NTP setup of our paper. The second component, which is only non-zero if there exists a target-context pair  $(z, j)$  with zero co-occurrence in the training set ( $\Pr(j, z) = 0$ ), specifies the directional component, akin to the role of  $L^{\text{mm}}$  in NTP. This component is absent in the analysis of [Levy & Goldberg \(2014\)](#), as they implicitly assume no zero entries in the co-occurrence matrix when setting the derivative of the loss to zero. However, when the co-occurrence matrix is sparse, zero derivation of the loss is not achieved at any finite solution. Instead,  $L_2$  identifies the infinite direction that drives the loss to its lower-bound. [Levy & Goldberg \(2014\)](#) address sparsity indirectly by discussing alternative matrix factorizations such as shifted positive-PMI (PPMI).

## B Proofs

### B.1 Centering of word embeddings

**Lemma 2.** For any  $\lambda > 0$ , any minimizer  $\widehat{\mathbf{W}}$  of *NTP-UFM* satisfies  $\mathbb{1}_V^\top \widehat{\mathbf{W}} = 0$ . That is, word embeddings are centered.

*Proof.* The proof is the same as [Thrapoulidis et al. \(2022, Lem. E.1\)](#) and is thus omitted.  $\square$

There is two consequences of Lemma 2:

- A. The word embedding matrix  $\mathbf{W}$  is centered motivating the use of the *centered* support matrix  $\tilde{\mathbf{S}}$  (rather than simply  $\mathbf{S}$ ) in Proxy **(P)**.
- B. The constraint  $\mathbb{1}^\top \mathbf{W} = 0$  can be added to *NTP-UFM* without changing its optimal set. Thus, the non-convex rank constraint can be ignored, making the problem convex, provided  $d \geq V - 1$ . Finally, note this also implies the constraint  $\mathbb{1}^\top \mathbf{L} = 0$  can be added without any change in Eq. (8).

### B.2 Proof of Lemma 1

We prove a slightly stronger version of Lemma 1 which assumes the additional constraint  $\mathbb{1}^\top \mathbf{L} = 0$  in Eq. *NTP-UFM*; see Lemma 2. The proof follows identically if we ignore this centering property.

As mentioned in the main body, the proof of the lemma essentially relies on the well-known fact that

$$\|\mathbf{L}\|_* = \min_{\mathbf{W}\mathbf{H}=\mathbf{L}} \frac{1}{2}\|\mathbf{W}\|^2 + \frac{1}{2}\|\mathbf{H}\|^2.$$

Firstly, fix any  $\mathbf{L}$  such that  $r := \text{rank}(\mathbf{L}) \leq d$  and  $\mathbb{1}_V^\top \mathbf{L} = 0$ . Let  $\mathbf{L} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top$  be its SVD with  $\mathbf{\Sigma} \in \mathbb{R}^{r \times r}$ . Choose

$$\widehat{\mathbf{W}} = \mathbf{U}\mathbf{\Sigma}^{1/2}\mathbf{R} \quad \text{and} \quad \widehat{\mathbf{H}} = \mathbf{R}^\top \mathbf{\Sigma}^{1/2}\mathbf{V}^\top$$

for some partial orthonormal matrix  $\mathbf{R} \in \mathbb{R}^{r \times d}$  (i.e.,  $\mathbf{R}^\top \mathbf{R} = \mathbb{I}_r$ ). It then holds clearly that

$$\widehat{\mathbf{W}}\widehat{\mathbf{H}} = \mathbf{L} \quad \text{and} \quad \|\widehat{\mathbf{W}}\|^2 + \|\widehat{\mathbf{H}}\|^2 = \text{tr}(\mathbf{U}\mathbf{\Sigma}\mathbf{U}^\top) + \text{tr}(\mathbf{V}\mathbf{\Sigma}\mathbf{V}^\top) = 2\text{tr}(\mathbf{\Sigma}) = 2\|\mathbf{L}\|_*,$$



Thus,

$$\text{CE}(\mathbf{L}) + \lambda \|\mathbf{L}\|_* = \text{CE}(\hat{\mathbf{W}}\hat{\mathbf{H}}) + \frac{\lambda}{2} \left( \|\hat{\mathbf{W}}\|^2 + \|\hat{\mathbf{H}}\|^2 \right).$$

Because we can apply this for any  $\mathbf{L}$ , it shows that<sup>6</sup>

$$\min_{\substack{\mathbf{L} \in \mathcal{R}_d \\ \mathbf{1}^\top \mathbf{L} = 0}} \text{CE}(\mathbf{L}) + \lambda \|\mathbf{L}\|_* \geq \min_{\mathbf{W}, \mathbf{H}} \text{CE}(\mathbf{W}\mathbf{H}) + \frac{\lambda}{2} \left( \|\mathbf{W}\|^2 + \|\mathbf{H}\|^2 \right), \quad (14)$$

where we define  $\mathcal{R}_d := \{\mathbf{L} \in \mathbb{R}^{V \times m} : \text{rank}(\mathbf{L}) \leq d\}$  to be the manifold of rank-constrained matrices.

Secondly, consider any  $\mathbf{W} \in \mathbb{R}^{V \times d}, \mathbf{H} \in \mathbb{R}^{d \times m}$  such that  $\mathbf{1}_V^\top \mathbf{W} = 0$ . Let  $\mathbf{L} = \mathbf{W}\mathbf{H}$  and denote  $\mathbf{L} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top$  its SVD. We then have  $\mathbf{1}^\top \mathbf{L} = 0$  and

$$\begin{aligned} \text{CE}(\mathbf{L}) + \lambda \|\mathbf{L}\|_* &= \text{CE}(\mathbf{W}\mathbf{H}) + \lambda \text{tr}(\mathbf{U}^\top \mathbf{W}\mathbf{H}\mathbf{V}) \\ &\leq \text{CE}(\mathbf{W}\mathbf{H}) + \frac{\lambda}{2} \left( \|\mathbf{U}^\top \mathbf{W}\|^2 + \|\mathbf{H}\mathbf{V}\|^2 \right) \\ &\leq \text{CE}(\mathbf{W}\mathbf{H}) + \frac{\lambda}{2} \left( \|\mathbf{W}\|^2 + \|\mathbf{H}\|^2 \right), \end{aligned} \quad (15)$$

where in the first line we used  $\|\mathbf{L}\|_* = \text{tr}(\mathbf{\Sigma}) = \text{tr}(\mathbf{U}^\top \mathbf{L}\mathbf{V})$ , in the second line we used Cauchy-Schwartz, and, in the third line we used that  $\|\mathbf{U}\|_2 \leq 1, \|\mathbf{V}\|_2 \leq 1$ . Thus, this proves

$$\begin{aligned} \min_{\mathbf{W}, \mathbf{H}} \text{CE}(\mathbf{W}\mathbf{H}) + \frac{\lambda}{2} \left( \|\mathbf{W}\|^2 + \|\mathbf{H}\|^2 \right) &= \min_{\substack{\mathbf{W}, \mathbf{H} \\ \mathbf{1}^\top \mathbf{W} = 0}} \text{CE}(\mathbf{W}\mathbf{H}) + \frac{\lambda}{2} \left( \|\mathbf{W}\|^2 + \|\mathbf{H}\|^2 \right) \\ &\geq \min_{\substack{\mathbf{L} \in \mathcal{R}_d \\ \mathbf{1}^\top \mathbf{L} = 0}} \text{CE}(\mathbf{L}) + \lambda \|\mathbf{L}\|_*, \end{aligned} \quad (16)$$

where the first equality follows by Lemma 2. Moreover, note that equalities in (15) hold if and only if the following three conditions hold

$$\begin{aligned} \mathbf{W}\mathbf{H} = \mathbf{L} = \mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top &\implies \mathbf{U}^\top \mathbf{W}\mathbf{H}\mathbf{V} = \mathbf{\Sigma} \\ \mathbf{U}^\top \mathbf{W} &= \mathbf{V}^\top \mathbf{H}^\top \\ \|\mathbf{U}^\top \mathbf{W}\|^2 &= \|\mathbf{W}\|^2, \|\mathbf{V}^\top \mathbf{H}^\top\|^2 = \|\mathbf{H}\|^2. \end{aligned}$$

The first two conditions combined give

$$\mathbf{U}^\top \mathbf{W}\mathbf{W}^\top \mathbf{U} = \mathbf{V}^\top \mathbf{H}^\top \mathbf{H}\mathbf{V} = \mathbf{\Sigma}.$$

Thus,  $\mathbf{W}\mathbf{W}^\top = \mathbf{U}\mathbf{\Sigma}\mathbf{U}^\top + \mathbf{U}_\perp \mathbf{\Sigma}_\perp \mathbf{U}_\perp^\top$  for some non-negative diagonal matrix  $\mathbf{\Sigma}_\perp$  and  $\mathbf{U}_\perp$  the orthogonal complement of  $\mathbf{U}$ . But, the third condition requires

$$\text{tr}(\mathbf{U}^\top \mathbf{W}\mathbf{W}^\top \mathbf{U}) = \text{tr}(\mathbf{W}\mathbf{W}^\top) \implies \text{tr}(\mathbf{\Sigma}) = \text{tr}(\mathbf{\Sigma}) + \text{tr}(\mathbf{\Sigma}_\perp) \implies \text{tr}(\mathbf{\Sigma}_\perp) = 0 \implies \mathbf{\Sigma}_\perp = 0.$$

This gives

$$\mathbf{W}\mathbf{W}^\top = \mathbf{U}\mathbf{\Sigma}\mathbf{U}^\top, \quad (17)$$

and similarly

$$\mathbf{H}^\top \mathbf{H} = \mathbf{V}\mathbf{\Sigma}\mathbf{V}^\top. \quad (18)$$

Combining Eqns. (14) and (16) proves that

$$\min_{\mathbf{W}, \mathbf{H}} \text{CE}(\mathbf{W}\mathbf{H}) + \frac{\lambda}{2} \left( \|\mathbf{W}\|^2 + \|\mathbf{H}\|^2 \right) = \min_{\substack{\mathbf{L} \in \mathcal{R}_d \\ \mathbf{1}^\top \mathbf{L} = 0}} \text{CE}(\mathbf{L}) + \lambda \|\mathbf{L}\|_*.$$

Moreover, the requirements for equalities in Eq. (15) established in Eqns. (17) and (18) show the desired about the form of gram matrices  $\mathbf{W}_\lambda \mathbf{W}_\lambda^\top$  and  $\mathbf{H}_\lambda \mathbf{H}_\lambda^\top$  as stated in the lemma.

<sup>6</sup>Here, use the fact that the minimizer of the LHS of (14) is attained because of coercivity of the nuclear norm.



### B.3 Proof of Theorems 1 and 2

The proof involves several components, and so, we proceed in steps.

#### B.3.1 Notations recap

We begin by recapping the essential notation used throughout this section.

Recall **NTP-SVM\***: The constraints in Eq. (9) require in-support logits to be equal to each other for every context. For fixed  $j \in [m]$ , the set of equality constraints for all  $z \neq z' \in \mathcal{S}_j$  is equivalent to the set of the same constraints for any anchor  $z_j \in \mathcal{S}_j$  and  $z' \neq z_j \in \mathcal{S}_j$ . That is, there is an effective total of  $S_j - 1$  linearly independent constraints for each  $j \in [m]$ .

**Definition 4** ( $E_{\text{in},j}$ ). For each  $j \in [m]$  fix an anchor index  $z_j \in \mathcal{S}_j$  and define matrix  $E_{\text{in},j} \in \mathbb{R}^{(S_j-1) \times V}$  with  $S_j - 1$  independent rows  $(e_{z_j} - e_z)^\top, z \neq z_j, z \in \mathcal{S}_j$ .

With this definition, we can rewrite Eq. (9) since

$$E_{\text{in},j} \ell_j = \mathbf{0}, \forall j \in [m] \Leftrightarrow \mathbf{L} = [\ell_1, \dots, \ell_m] \text{ satisfies Eq. (9)}$$

The inequality constraints in Eq. (10), impose the requirement that in-support logits must be greater than out-of-support logits. Considering Eq. (9), we can again select any anchor  $z_j \in \mathcal{S}_j$  and reformulate Eq. (10) as  $E_{\text{out},j} \ell_j \geq \mathbf{1}, j \in [m]$ . Here,  $E_{\text{out},j} \in \mathbb{R}^{(V-S_j) \times V}$  has  $V - S_j$  independent rows  $(e_{z_j} - e_v)^\top, v \notin \mathcal{S}_j$ . In total, there are  $V - 1$  constraints per context in **NTP-SVM\***.

With the above definition of the matrices  $E_{\text{in},j} \in \mathbb{R}^{(S_j-1) \times V}, j \in [m]$ , we can define the subspace  $\mathcal{F}$  of  $V \times m$  matrices that have their  $j$ -th column living on the range space  $\mathcal{R}(E_{\text{in},j}^\top)$  for all  $j \in [m]$ :

$$\mathcal{F} := \left\{ \mathbf{L} \in \mathbb{R}^{V \times m} : \ell_j \in \mathcal{R}(E_{\text{in},j}^\top), j \in [m] \right\} = \text{span} \left( \left\{ (e_z - e_{z'})^\top \tilde{e}_j : z \neq z' \in \mathcal{S}_j, j \in [m] \right\} \right).$$

Note that this set depends on the support matrix  $\mathbf{S}$ , but we suppress the dependence to simplify notation. Let  $\mathcal{P}_{\mathcal{F}} : \mathbb{R}^{V \times m} \rightarrow \mathbb{R}^{V \times m}$  be the projection operator on this subspace. For concreteness, note that this can be explicitly defined as the matrix operator projecting each column  $\ell_j$  of  $\mathbf{L}$  onto  $\mathcal{R}(E_{\text{in},j}^\top)$ , i.e.,

$$\mathcal{P}_{\mathcal{F}}(\mathbf{L}) = [\mathcal{P}_{\mathcal{F}}(\mathbf{L})_1 \quad \mathcal{P}_{\mathcal{F}}(\mathbf{L})_2 \quad \dots \quad \mathcal{P}_{\mathcal{F}}(\mathbf{L})_m] \text{ with } \mathcal{P}_{\mathcal{F}}(\mathbf{L})_j := E_{\text{in},j}^\top (E_{\text{in},j} E_{\text{in},j}^\top)^{-1} E_{\text{in},j} \ell_j, j \in [m].$$

Note that  $\mathcal{P}_{\mathcal{F}}$  is a projection operator since it is clearly linear and  $\mathcal{P}_{\mathcal{F}}(\mathcal{P}_{\mathcal{F}}(\mathbf{L})) = \mathcal{P}_{\mathcal{F}}(\mathbf{L})$ . Also, it is clear that  $\mathcal{P}_{\mathcal{F}}(\mathbf{L}) \in \mathcal{F}$  for all  $\mathbf{L} \in \mathbb{R}^{V \times m}$ . Operationally, note that  $\mathcal{P}_{\mathcal{F}}(\mathbf{L})$  has zero-entries on off-support elements. We will use this property indiscriminately throughout this section.

Finally, define the orthogonal complement to  $\mathcal{F}$  as

$$\mathcal{F}_\perp := \left\{ \mathbf{L} \in \mathbb{R}^{V \times m} : \ell_j \in \mathcal{N}(E_{\text{in},j}), j \in [m] \right\} = \left\{ \mathbf{L} \in \mathbb{R}^{V \times m} : E_{\text{in},j} \ell_j = \mathbf{0}, j \in [m] \right\}. \quad (19)$$

Clearly,  $\mathbb{R}^{V \times m} = \mathcal{F} \oplus \mathcal{F}_\perp$ . Operationally, note that  $\mathbf{L} \in \mathcal{F}_\perp$  implies that on-support elements of every column are identical, that is,  $\mathbf{L}[z, j] = \mathbf{L}[z', j]$  for all  $z \in \mathcal{S}_j$  and  $j \in [m]$ . In other words, it holds

$$\mathbf{L} \in \mathcal{F}_\perp \iff \mathbf{L} \text{ satisfies (9)}, \quad (20)$$

which we will use without further explicit reference.

We also let  $\mathcal{P}_\perp$  denote the projection operator to  $\mathcal{F}_\perp$ .

### B.3.2 When does the NTP loss attain its entropy lower-bound?

For word/context embeddings  $\mathbf{W}, \mathbf{H}$  and logits  $\mathbf{L} = \mathbf{WH}$ , recall the NTP loss is written as

$$\text{CE}(\mathbf{WH}) = \text{CE}(\mathbf{L}) = \sum_{j \in [m]} \hat{\pi}_j \sum_{z \in \mathcal{S}_j} \hat{p}_{j,z} \log \left( 1 + \sum_{z' \neq z \in \mathcal{S}_j} e^{-(L[z,j] - L[z',j])} + \sum_{v \notin \mathcal{S}_j} e^{-(L[z,j] - L[v,j])} \right), \quad (21)$$

and is lower bounded by the empirical  $T$ -gram entropy of the data (Shannon, 1948), i.e., for all  $\theta'$ :  $\text{CE}(\mathbf{WH}) \geq \mathcal{H} := \hat{\mathbb{E}}_{(x,z) \sim \mathcal{T}_n} [-\log(\hat{p}(z|x))]$ .

We now state explicit conditions on the logit matrix  $\mathbf{L}$  for which the lower bound is attained by a softmax model  $\mathbf{L} = \mathbf{S}(\mathbf{WH})$  with unconstrained context embeddings  $\mathbf{H}$ . Due to sparsity of the conditional probability matrix  $\mathbf{P}$ , the lower bound is *not* attained for any finite logit matrix, but it can be reached asymptotically provided the conditions stated in the proposition below hold.

The statement below is a slight extension of Thrampoulidis (2024, Prop. 1).

**Proposition 4.** *The NTP objective of a softmax model  $\mathbf{S}(\mathbf{WH})$  with  $\mathbf{W} \in \mathbb{R}^{V \times d}$ ,  $\mathbf{H} \in \mathbb{R}^{d \times m}$  can reach the empirical entropy lower-bound if the following three conditions hold simultaneously*

A. [NTP- $\mathcal{H}$ -compatibility] *There exists logit matrix  $\mathbf{L}_1 \in \mathbb{R}^{V \times m}$  such that*

$$\forall j \in [m], z \neq z' \in \mathcal{S}_j : \mathbf{L}_1[j, z] - \mathbf{L}_1[j, z'] = \log \left( \frac{\hat{p}_{j,z}}{\hat{p}_{j,z'}} \right). \quad (22)$$

B. [NTP-separability] *There exists logit matrix  $\mathbf{L}_2 \in \mathbb{R}^{V \times m}$  that is feasible in NTP-SVM $_{\star}$ .*

C. [Rank constraint] *For all arbitrarily large positive scalars  $\rho$ ,*

$$\text{rank}(\mathbf{L}_1 + \rho \cdot \mathbf{L}_2) \leq d.$$

Concretely, if the above conditions hold, then for all arbitrarily large  $\rho$ , there exist  $\mathbf{W}_\rho \in \mathbb{R}^{V \times d}$  and  $\mathbf{H}_\rho \in \mathbb{R}^{d \times m}$  such that

$$\lim_{\rho \rightarrow \infty} \text{CE}(\mathbf{W}_\rho \mathbf{H}_\rho) = \mathcal{H}.$$

*Proof.* It suffices to prove that for  $\mathbf{L}_\rho := \mathbf{L}_1 + \rho \cdot \mathbf{L}_2$ :

$$\lim_{\rho \rightarrow \infty} \text{CE}(\mathbf{L}_\rho) = \mathcal{H}. \quad (23)$$

To see that this is sufficient, note from condition C of the theorem that for all arbitrarily large  $\rho$ ,  $\mathbf{L}_\rho$  admits a factorization  $\mathbf{L}_\rho = \mathbf{W}_\rho \mathbf{H}_\rho$  where  $\mathbf{W}_\rho \in \mathbb{R}^{V \times d}$  and  $\mathbf{H}_\rho \in \mathbb{R}^{d \times m}$ . Since this holds for all large  $\rho$ , Eq. (23) implies the desired.

Thus, we now prove Eq. (23). For any  $j \in [m]$ , it is easy to check, using conditions A and B of the theorem, that the following two statements are true for all  $\rho > 0$ :

A. For all  $z \neq z' \in \mathcal{S}_j$ :

$$\begin{aligned} \mathbf{L}_\rho[z, j] - \mathbf{L}_\rho[z', j] &= (\mathbf{L}_1[z, j] - \mathbf{L}_1[z', j]) + \rho \cdot (\mathbf{L}_2[z, j] - \mathbf{L}_2[z', j]) \\ &= \log(\hat{p}_{j,z} / \hat{p}_{j,z'}) \end{aligned}$$

B. For all  $z \in \mathcal{S}_j, v \notin \mathcal{S}_j$ :

$$\begin{aligned} \mathbf{L}_\rho[z, j] - \mathbf{L}_\rho[v, j] &= (\mathbf{L}_1[z, j] - \mathbf{L}_1[v, j]) + \rho \cdot (\mathbf{L}_2[z, j] - \mathbf{L}_2[v, j]) \\ &\geq (\mathbf{L}_1[z, j] - \mathbf{L}_1[v, j]) + \rho \\ &\geq -2\|\mathbf{L}_1\|_2 + \rho. \end{aligned}$$

In the last inequality above, we used the loose bound  $\|\mathbf{L}_1\|_2 \geq \max_{j \in [m], z \in [V]} |\mathbf{L}_1[z, j]|$  for the spectral norm of  $\mathbf{L}_1$ .

Using these in Eq. (21) gives:

$$\begin{aligned}
\text{CE}(\mathbf{L}_\rho) &\leq \sum_{j \in [m]} \hat{\pi}_j \sum_{z \in \mathcal{S}_j} \hat{p}_{j,z} \log \left( 1 + \sum_{z' \neq z \in \mathcal{S}_j} \frac{\hat{p}_{j,z'}}{\hat{p}_{j,z}} + Ve^{-\rho} e^{-2\|\mathbf{L}_1\|_2} \right) \\
&= \sum_{j \in [m]} \hat{\pi}_j \sum_{z \in \mathcal{S}_j} \hat{p}_{j,z} \log \left( \frac{1}{\hat{p}_{j,z}} + Ve^{-\rho} e^{-2\|\mathbf{L}_1\|_2} \right) \\
&= \sum_{j \in [m]} \hat{\pi}_j \sum_{z \in \mathcal{S}_j} \hat{p}_{j,z} \log \left( \frac{1}{\hat{p}_{j,z}} \right) + \sum_{j \in [m]} \hat{\pi}_j \sum_{z \in \mathcal{S}_j} \hat{p}_{j,z} \log \left( 1 + \hat{p}_{j,z} Ve^{-\rho} e^{-2\|\mathbf{L}_1\|_2} \right). \\
&\leq \mathcal{H} + Ve^{-2\|\mathbf{L}_1\|_2} e^{-\rho}. \tag{24}
\end{aligned}$$

The last inequality uses  $\log(1+x) \leq x, x > 0$ .

Clearly the bound above converges to  $\mathcal{H} = \sum_{j \in [m]} \hat{\pi}_j \sum_{z \in \mathcal{S}_j} \hat{p}_{j,z} \log \left( \frac{1}{\hat{p}_{j,z}} \right)$  as  $\rho \rightarrow \infty$ . From this and  $\text{CE}(\mathbf{L}_\rho) \geq \mathcal{H}$ , the desired Eq. (23) follows. This completes the proof.  $\square$

Note that the first two conditions of the theorem are easy to satisfy.

On the one hand, Eq. (22) always has a solution. For later use, we also note that the solution is unique on the subspace  $\mathcal{F}$  and we denote this solution as  $\mathbf{L}^{\text{in}}$ . This is easy to check: for any  $j \in [m]$ ,  $\ell_j$  is a solution of Eq. (22) if and only if it takes the form  $\ell_j = \mathbf{E}_{\text{in},j}^\top \left( \mathbf{E}_{\text{in},j} \mathbf{E}_{\text{in},j}^\top \right)^{-1} \mathbf{a}_j + \mathbf{v}$ ,  $\mathbf{v} \in \mathcal{N}(\mathbf{E}_{\text{in}})$  where  $\mathbf{a}_j \in \mathbb{R}^{\mathcal{S}_j-1}$  has entries  $\log(\hat{p}_{j,z}/\hat{p}_{j,z'})$ ,  $z' \neq z \in \mathcal{S}_j$ . Thus,  $\mathbf{L}^{\text{in}}$  is explicitly defined as follows.

**Lemma 3** ( $\mathbf{L}^{\text{in}}$ ). *Let  $\mathbf{L}^{\text{in}}$  be the unique solution of Eq. (22) on the subspace  $\mathcal{F}$ . Then,  $\mathbf{L}^{\text{in}}$  has columns*

$$\ell_j^{\text{in}} = \mathbf{E}_{\text{in},j}^\top \left( \mathbf{E}_{\text{in},j} \mathbf{E}_{\text{in},j}^\top \right)^{-1} \mathbf{a}_j, \quad \forall j \in [m],$$

where  $\mathbf{a}_j \in \mathbb{R}^{\mathcal{S}_j-1}$  has entries  $\log(\hat{p}_{j,z_j}/\hat{p}_{j,z'})$  for  $z' \neq z_j \in \mathcal{S}_j$ .

On the other hand, as mentioned in Sec. 4.1, **NTP-SVM $\star$**  is always feasible.

The challenge is in satisfying conditions A-C simultaneously.

However, when  $d \geq V$ , then the rank condition C of the theorem is trivial. Thus, we have the following corollary.

**Corollary 2.** *Suppose  $d \geq V$ , then the entropy lower bound can be asymptotically attained by the NTP loss.*

### B.3.3 Ball-constrained CE minimization and main result

Instead of the regularized problem (8), it is more convenient (but, equivalent due to convexity) to study the following ball-constrained version:

$$\hat{\mathbf{L}}_B \in \arg \min_{\|\mathbf{L}\|_* \leq B} \text{CE}(\mathbf{L}), \tag{25}$$

parameterized by  $B > 0$ . We will study properties of the minimizers  $\hat{\mathbf{L}}_B$  of (25) as  $B \rightarrow \infty$ . The theorem below subsumes Theorems 1 and 2.

**Theorem 3.** *Suppose  $d \geq V$ . Denote  $\hat{\mathbf{L}}_B$  any solution of (25) for regularization  $B$ . Then, in the limit of diverging regularization the following statements are true.*

- (i)  $\lim_{B \rightarrow \infty} \text{CE}(\hat{\mathbf{L}}_B) = \mathcal{H}$ .
- (ii)  $\lim_{B \rightarrow \infty} \mathcal{P}_{\mathcal{F}}(\hat{\mathbf{L}}_B) = \mathbf{L}^{\text{in}}$ , where  $\mathbf{L}^{\text{in}}$  is defined in Lemma 3.

(iii) The solution diverges in norm:  $\lim_{B \rightarrow \infty} \|\widehat{\mathbf{L}}_B\| = +\infty$ .

(iv) The solution converges in direction to the optimal set of  $\text{NTP-SVM}_*$ . That is, there exists minimizer  $\mathbf{L}^{\text{mm}}$  of  $\text{NTP-SVM}_*$  such that

$$\lim_{B \rightarrow \infty} \left\| \frac{\widehat{\mathbf{L}}_B}{\|\widehat{\mathbf{L}}_B\|_*} - \frac{\mathbf{L}^{\text{mm}}}{\|\mathbf{L}^{\text{mm}}\|_*} \right\| = 0.$$

### B.3.4 The in-support NTP loss and its minimizer $\mathbf{L}^{\text{in}}$

Define the in-support NTP loss as follows:

$$\text{CE}_{\text{in}}(\mathbf{L}) = \sum_{j \in [m]} \hat{\pi}_j \sum_{z \in \mathcal{S}_j} \hat{p}_{j,z} \log \left( 1 + \sum_{z' \neq z \in \mathcal{S}_j} e^{-(L[z,j] - L[z',j])} \right)$$

Contrary to (21), the summation inside the logarithm is only over word indices  $z$  that belong to the support set of context  $j$ .

By non-negativity of log, for all  $\mathbf{L}$  it holds  $\text{CE}(\mathbf{L}) \geq \text{CE}_{\text{in}}(\mathbf{L})$ . In fact, because of the sparsity ansatz ( $\exists j \in [m]$  such that  $S_j < V$ ) and strict nonnegativity of exponential, the inequality is strict for all finite  $\mathbf{L}$ , i.e.

$$\text{CE}(\mathbf{L}) > \text{CE}_{\text{in}}(\mathbf{L}), \forall \mathbf{L} \in \mathbb{R}^{V \times d}. \quad (26)$$

Moreover, observe that for any  $\mathbf{L}$ , it holds that

$$\text{CE}_{\text{in}}(\mathbf{L}) = \text{CE}_{\text{in}}(\mathcal{P}_{\mathcal{F}}(\mathbf{L})). \quad (27)$$

To see this, decompose  $\mathbf{L}$  in orthogonal components onto  $\mathcal{F}$  and  $\mathcal{F}_{\perp}$ , respectively, i.e.,  $\mathbf{L} = \mathcal{P}_{\mathcal{F}}(\mathbf{L}) + \mathcal{P}_{\perp}(\mathbf{L})$ . By Eq. (20),  $\mathbf{L}_{\perp} := \mathcal{P}_{\perp}(\mathbf{L})$  satisfies Eq. (9), i.e.  $\mathbf{L}_{\perp}[z,j] - \mathbf{L}_{\perp}[z',j] = 0$  for all  $z, z' \in \mathcal{S}_j$  and  $j \in [m]$ .

The lemma below uses this to show that  $\text{CE}_{\text{in}}$  has a unique minimizer in  $\mathcal{F}$ .

**Lemma 4.** *The in-support loss  $\text{CE}_{\text{in}}(\mathbf{L})$  has a unique finite minimizer in  $\mathcal{F}$ , which we denote  $\mathbf{L}^{\text{in}}$ :*

$$\mathbf{L}^{\text{in}} = \arg \min_{\mathbf{L} \in \mathcal{F}} \text{CE}_{\text{in}}(\mathbf{L}). \quad (28)$$

The minimizer  $\mathbf{L}^{\text{in}}$  satisfies the NTP $_{\mathcal{H}}$ -compatibility condition in Eq. (22) and

$$\text{CE}_{\text{in},*} := \text{CE}_{\text{in}}(\mathbf{L}^{\text{in}}) = \mathcal{H}.$$

*Proof.* Note that  $\mathbf{L} \in \mathcal{F}$  is a column-wise separable constraint  $\ell_j \in \mathcal{R}(E_{\text{in},j}^{\top})$  that restricts the non-zero entries  $L[z,j] = \ell_{j,z}$  to indices  $z \in \mathcal{S}_j$ . The objective is also separable with respect to columns. Without loss of generality fix  $j = 1$  and assume  $\mathcal{S}_j = [S_1]$ . It then suffices proving that the function  $f : \mathbb{R}^{S_1} \rightarrow \mathbb{R}$ :

$$f(\ell) = - \sum_{z \in S_1} \hat{p}_{1,z} \log \left( \frac{\exp(\ell_z)}{\sum_{z' \in S_1} \exp(\ell_{z'})} \right), \quad (29)$$

has a unique minimizer in  $\mathcal{R}(\widetilde{E}_{\text{in},1}^{\top})$ , where  $\widetilde{E}_{\text{in},1} := E_{\text{in},1}[\cdot, S_1] \in \mathbb{R}^{(S_1-1) \times S_1}$  is the non-zero block of  $E_{\text{in},1}$ .

By non-negativity of the KL divergence,

$$f(\ell) \geq - \sum_{z \in S_1} \hat{p}_{1,z} \log(\hat{p}_{1,z})$$

with equality if and only if,

$$\forall z \in S_1 : \frac{\exp(\ell_z)}{\sum_{z' \in S_1} \exp(\ell_{z'})} = \hat{p}_{1,z}. \quad (30)$$

In turn, this is equivalent to

$$\forall z, z' \in \mathcal{S}_1 : \ell_z - \ell_{z'} = \log(\hat{p}_{1,z} / \hat{p}_{1,z'}). \quad (31)$$

To see the equivalence note the following: On the one hand, starting with Eq. (30) for  $z, z' \in \mathcal{S}_1$ , dividing both sides of the two equalities, and taking the logarithm, we arrive at Eq. (31). On the other hand, if Eq. (31) holds, then Eq. (30) holds since

$$\frac{\sum_{z' \in \mathcal{S}_1} \exp(\ell_{z'})}{\exp(\ell_z)} = \sum_{z' \in \mathcal{S}_1} \exp(\ell_{z'} - \ell_z) = \frac{\sum_{z' \in \mathcal{S}_1} \hat{p}_{1,z'}}{\hat{p}_{1,z}} = \frac{1}{\hat{p}_{1,z}}.$$

Now, observe that (31) is the same as the NTP $_{\mathcal{H}}$ -compatibility equation for  $j = 1$ . Moreover, note that a logit vector  $\ell \in \mathbb{R}^{\mathcal{S}_1}$  solves (31) if and only if for

$$\ell = \tilde{\mathbf{E}}_{\text{in},1}^{\top} \left( \tilde{\mathbf{E}}_{\text{in},1} \tilde{\mathbf{E}}_{\text{in},1}^{\top} \right)^{-1} \mathbf{a}_1 + \mathbf{v}, \quad \mathbf{v} \in \mathcal{N} \left( \tilde{\mathbf{E}}_{\text{in},1} \right),$$

where  $\mathbf{a} \in \mathbb{R}^{\mathcal{S}_1-1}$  with entries  $\log(\hat{p}_{1,z_1} / \hat{p}_{1,z'})$  for  $z' \neq z_1 \in \mathcal{S}_1$ . This proves that there is a unique minimizer  $\ell_1^{\text{in}} := \tilde{\mathbf{E}}_{\text{in},1}^{\top} \left( \tilde{\mathbf{E}}_{\text{in},1} \tilde{\mathbf{E}}_{\text{in},1}^{\top} \right)^{-1} \mathbf{a}_1$  on  $\mathcal{R} \left( \tilde{\mathbf{E}}_{\text{in},1} \right)$ .

Since the choice of  $j = 1$  above was arbitrary, this proves the lemma.  $\square$

### B.3.5 Auxiliary lemmas

This section proves a sequence of auxiliary lemmas that are used to prove Thm. 3 in the next section.

The first lemma shows that entropy can only be approached asymptotically, i.e., provided the logit matrix diverges.

**Lemma 5.** For all  $\mathbf{L} \in \mathbb{R}^{V \times m}$  and all  $\mathbf{L}' \in \mathcal{F}_{\perp}$  that additionally satisfy  $\mathbf{E}_{\text{out},j} \ell_j' \geq 1, \forall j \in [m]$  (i.e.,  $\mathbf{L}'$  is NTP-SVM $_{\star}$  feasible), it holds

$$\text{CE}_{\text{in}}(\mathbf{L}) = \lim_{R \rightarrow \infty} \text{CE}(\mathbf{L} + R\mathbf{L}').$$

Using this, it holds for all  $\mathbf{L} \in \mathbb{R}^{V \times m}$  that

$$\lim_{R \rightarrow \infty} \text{CE}(\mathbf{L}^{\text{in}} + R\mathbf{L}') = \text{CE}_{\text{in}}(\mathbf{L}^{\text{in}}) = \text{CE}_{\text{in},\star} = \mathcal{H} \leq \text{CE}_{\text{in}}(\mathbf{L}) < \text{CE}(\mathbf{L}). \quad (32)$$

*Proof.* To see the first claim note that  $\mathbf{L}' \in \mathcal{F}_{\perp}$  implies

$$\begin{aligned} \forall j \in [m], z \neq z' \in \mathcal{S}_j : \mathbf{L}'[j, z] = \mathbf{L}'[j, z'] &\implies \\ \forall j \in [m], z \in \mathcal{S}_j : \sum_{z' \in \mathcal{S}_j} e^{-(\mathbf{L}[z,j] - \mathbf{L}[z',j] + R(\mathbf{L}'[j,z] - \mathbf{L}'[j,z']))} &= \sum_{z' \in \mathcal{S}_j} e^{-(\mathbf{L}[z,j] - \mathbf{L}[z',j])} \end{aligned}$$

and  $\mathbf{E}_{\text{out},j} \ell_j' \geq 1, \forall j \in [m]$  implies

$$\begin{aligned} \forall j \in [m], z \in \mathcal{S}_j, v \notin \mathcal{S}_j : \mathbf{L}'[j, z] - \mathbf{L}'[j, v] > 0 &\implies \\ \forall j \in [m], z \in \mathcal{S}_j : \sum_{v \notin \mathcal{S}_j} e^{-(\mathbf{L}[z,j] - \mathbf{L}[v,j] + R(\mathbf{L}'[j,z] - \mathbf{L}'[j,v]))} &\xrightarrow{R \rightarrow \infty} 0. \end{aligned}$$

The second claim applies the first claim for  $\mathbf{L} = \mathbf{L}^{\text{in}}$  and uses Eq. (26).  $\square$

We now use the previous lemma to show that the ball-constrained minimizer diverges as the ball constraint approaches infinity.

**Lemma 6.** The norm of the ball-constrained minimizer diverges, that is

$$\lim_{B \rightarrow \infty} \|\hat{\mathbf{L}}_B\|_{\star} = \infty$$

*Proof.* Suppose on the contrary that  $\lim_{B \rightarrow \infty} \|\widehat{\mathbf{L}}_B\|_* < \infty$ . From Eq. (32), we have

$$\text{CE}(\widehat{\mathbf{L}}_B) > \lim_{R \rightarrow \infty} \text{CE}(\mathbf{L}^{\text{in}} + R\mathbf{L}^{\text{mm}}).$$

Moreover, by triangle inequality for nuclear norm  $\|\mathbf{L}_R\|_* \leq \|\mathbf{L}^{\text{in}}\|_* + R\|\mathbf{L}^{\text{mm}}\|_*$ . Thus, for any arbitrarily large  $R > 0$ , there exists large enough  $B$  such that  $\|\mathbf{L}_R\|_* \leq B$ , making  $\mathbf{L}_R := \mathbf{L}^{\text{in}} + R\mathbf{L}^{\text{mm}}$  feasible in (25). These together, contradict optimality of  $\widehat{\mathbf{L}}_B$ , proving that  $\lim_{B \rightarrow \infty} \|\widehat{\mathbf{L}}_B\|_* = \infty$ . Finally, recalling the norm inequality  $\|\widehat{\mathbf{L}}_B\| \geq \|\widehat{\mathbf{L}}_B\|_* / \sqrt{V}$ , it also holds that  $\lim_{B \rightarrow \infty} \|\widehat{\mathbf{L}}_B\| = \infty$ .  $\square$

The last lemma constructs a sequence of logits approaching the entropy lower bound while remaining feasible in problem (25) as  $B \rightarrow \infty$ .

**Lemma 7.** *The objective function approaches the loss lower bound, i.e.,*

$$\lim_{B \rightarrow \infty} \text{CE}(\widehat{\mathbf{L}}_B) = \text{CE}_{\text{in},*} = \mathcal{H}.$$

*Concretely, there exists logit matrix  $\mathbf{L}_R$  parameterized by  $R = R(B)$  that is feasible in (25) and satisfies:*

$$\text{CE}(\mathbf{L}_R) \leq \text{CE}_{\text{in},*} + Ve^{2\|\mathbf{L}^{\text{in}}\|_2} e^{-R}. \quad (33)$$

*Moreover,  $\lim_{B \rightarrow \infty} R = \infty$  and  $\lim_{B \rightarrow \infty} \frac{R}{\|\widehat{\mathbf{L}}_B\|_*} = \frac{1}{\|\mathbf{L}^{\text{mm}}\|_*}$  where  $\|\mathbf{L}^{\text{mm}}\|_*$  is the optimal cost of NTP-SVM $_*$ .*

*Proof.* We evaluate the loss achieved by the following candidate good point

$$\mathbf{L}_R := \mathbf{L}^{\text{in}} + R(B) \cdot \mathbf{L}^{\text{mm}}.$$

Here, (for large enough  $B > \|\mathbf{L}^{\text{in}}\|_*$ ) we set

$$R := R(B) = \frac{\|\widehat{\mathbf{L}}_B\|_* - \|\mathbf{L}^{\text{in}}\|_*}{\|\mathbf{L}^{\text{mm}}\|_*}.$$

Note this is chosen so that  $\|\mathbf{L}_R\|_* \leq \|\widehat{\mathbf{L}}_B\|_* \leq B$  towards making  $\mathbf{L}_R$  feasible in (25). Also, recall from Lemma 6 that  $\|\widehat{\mathbf{L}}_B\|_* \rightarrow \infty$  as  $B \rightarrow \infty$ ; thus, also  $R \rightarrow \infty$ .

It remains to prove Eq. (33). This follows directly from the sequence of equations in (24).  $\square$

### B.3.6 Proof of Theorem 3

We prove each statement separately.

**Statement (i).** Since for all  $\mathbf{L}$ , it holds  $\text{CE}(\mathbf{L}) \geq \mathcal{H}$ , it suffices to prove that  $\mathcal{H}$  can be asymptotically attained. Since  $d \geq V$ , this follows from Corollary 2. (For a more explicit proof, see Lemma 7).

**Statement (ii).** Since for all  $\mathbf{L}$  it holds  $\text{CE}(\mathbf{L}) \geq \text{CE}_{\text{in}}(\mathbf{L}) \geq \text{CE}_{\text{in},*} = \mathcal{H}$ , it follows from Statement (i) that

$$\lim_{B \rightarrow \infty} \text{CE}_{\text{in}}(\widehat{\mathbf{L}}_B) = \text{CE}_{\text{in},*}.$$

But  $\text{CE}_{\text{in}}(\mathbf{L}) = \text{CE}_{\text{in}}(\mathcal{P}_{\mathcal{F}}(\mathbf{L}))$  (see Eq. (27)). Thus,

$$\lim_{B \rightarrow \infty} \text{CE}_{\text{in}}(\mathcal{P}_{\mathcal{F}}(\widehat{\mathbf{L}}_B)) = \text{CE}_{\text{in},*}.$$

The desired now follows since  $\mathbf{L}^{\text{in}}$  is unique minimizer of  $\text{CE}_{\text{in}}$  on  $\mathcal{T}$  by Lemma 4.

**Statement (iii).** Follows directly from Lemma 6.

**Statement (iv).** From Lemma 7, there exists  $\mathbf{L}_R$  for which the loss is close to  $\mathcal{H}$  (see Eq. (33)).



Towards arriving at a contradiction, we will show that if  $\widehat{\mathbf{L}}_B$  is not in the direction of  $\mathbf{L}^{\text{mm}}$ , then it incurs a loss that is larger than the upper bound of  $\text{CE}(\mathbf{L}_R)$  computed in Eq. (33).

To do this, assuming the statement of the theorem is not true, we will lower bound

$$\text{CE}(\widehat{\mathbf{L}}_B) - \mathcal{H} = \sum_{j \in [m]} \hat{\pi}_j \sum_{z \in \mathcal{S}_j} \hat{p}_{j,z} \log \left( \hat{p}_{j,z} \left( \sum_{z' \in \mathcal{S}_j} e^{-(\widehat{\mathbf{L}}_B[j,z] - \widehat{\mathbf{L}}_B[j,z'])} + \sum_{v \notin \mathcal{S}_j} e^{-(\widehat{\mathbf{L}}_B[j,z] - \widehat{\mathbf{L}}_B[j,v])} \right) \right). \quad (34)$$

By our assumption, there exists  $\epsilon > 0$ , such that there exists arbitrarily large  $B$  satisfying:

$$\left\| \frac{\|\mathbf{L}^{\text{mm}}\|_*}{\|\widehat{\mathbf{L}}_B\|_*} \widehat{\mathbf{L}}_B - \mathbf{L}^{\text{mm}} \right\| > \epsilon, \quad (35)$$

for all minimizers  $\mathbf{L}^{\text{mm}}$  of **NTP-SVM\***. Define

$$\widehat{\mathbf{L}} = \frac{1}{R'(B)} (\widehat{\mathbf{L}}_B - \mathbf{L}^{\text{in}}),$$

where,  $R' := R'(B) > 0$  is chosen so that  $\|\widehat{\mathbf{L}}\|_* < \|\mathbf{L}^{\text{mm}}\|_*$ . Concretely, this can be ensured by setting (to see this use the triangle inequality for nuclear norm):

$$R' = \frac{(\|\widehat{\mathbf{L}}_B\|_* + \|\mathbf{L}^{\text{in}}\|_* + \zeta)}{\|\mathbf{L}^{\text{mm}}\|_*},$$

for some  $\zeta > 0$ . Recall that  $\|\widehat{\mathbf{L}}_B\|_* \rightarrow \infty$  as  $B \rightarrow \infty$ . On the other hand  $\|\mathbf{L}^{\text{in}}\|_*$  is not dependent on  $B$ . Thus, in the large limit  $B \rightarrow \infty$ , it holds  $\left| \frac{\|\mathbf{L}^{\text{mm}}\|_*}{\|\widehat{\mathbf{L}}_B\|_*} - \frac{1}{R'} \right| \rightarrow 0$  and  $\frac{\|\mathbf{L}^{\text{in}}\|_*}{\|\widehat{\mathbf{L}}_B\|_*} \rightarrow 0$ . These combined with (35) show that we can always choose  $B$  large enough so that Eq. (35) guarantees, for some  $\epsilon' > 0$ , that

$$\|\widehat{\mathbf{L}} - \mathbf{L}^{\text{mm}}\| \geq \epsilon'.$$

But  $\widehat{\mathbf{L}}$  achieves the optimal cost of **NTP-SVM\***, since  $\|\widehat{\mathbf{L}}\|_* < \|\mathbf{L}^{\text{mm}}\|_*$ . Thus, there exists  $\delta \in (0, 1)$  and  $j \in [m]$  such that at least one of the following is true

(i)  $\exists z$  and  $z' \neq z \in \mathcal{S}_j$  such that

$$|\widehat{\mathbf{L}}[j, z] - \widehat{\mathbf{L}}[j, z']| \geq \delta, \quad (36)$$

(ii)  $\exists z \in \mathcal{S}_j, v \notin \mathcal{S}_j$  such that

$$\widehat{\mathbf{L}}[j, z] - \widehat{\mathbf{L}}[j, v] \leq 1 - \delta. \quad (37)$$

Case (i): Without loss of generality  $(\mathbf{e}_z - \mathbf{e}_{z'})^\top \widehat{\boldsymbol{\ell}}_j \leq -\delta$  (otherwise, flip  $z, z'$ ). Thus, ignoring all but one term in (34) gives

$$\text{CE}(\widehat{\mathbf{L}}_B) - \mathcal{H} \geq \hat{\pi}_j \hat{p}_{j,z} \log \left( \hat{p}_{j,z} e^{-(\widehat{\mathbf{L}}_B[j,z] - \widehat{\mathbf{L}}_B[j,z'])} \right). \quad (38)$$

But,

$$\widehat{\mathbf{L}}_B[j, z] - \widehat{\mathbf{L}}_B[j, z'] = R'(\widehat{\mathbf{L}}[j, z] - \widehat{\mathbf{L}}[j, z']) + (\mathbf{L}^{\text{in}}[j, z] - \mathbf{L}^{\text{in}}[j, z']) \leq -R'\delta + 2\|\mathbf{L}^{\text{in}}\|_2 \quad (39)$$

Put in (34) and using  $\hat{p}_{j,z} \geq 1/V$ ,  $\hat{\pi}_j \geq 1/m$  shows

$$\text{CE}(\widehat{\mathbf{L}}_B) \geq \mathcal{H} + \frac{1}{mV} \log \left( \frac{1}{V} e^{R'\delta} \exp(-2\|\mathbf{L}^{\text{in}}\|_2) \right) \geq \mathcal{H} + \frac{1}{mV} \log \left( \frac{e^{R'\delta}}{V \exp(2\|\mathbf{L}^{\text{in}}\|_2)} \right)$$

Compare this with (33), it is clear that as  $B \rightarrow \infty$ , so that  $R' \rightarrow \infty$ , it holds  $\frac{1}{mV} \log \left( \frac{e^{R'\delta}}{V \exp(2\|\mathbf{L}^{\text{in}}\|_2)} \right) > 1 > V \exp(2\|\mathbf{L}^{\text{in}}\|_2) e^{-R}$ . Thus,  $\text{CE}(\widehat{\mathbf{L}}_B) > \text{CE}(\mathbf{L}_R^{\text{in}})$ , a contradiction.

*Case (ii):* We can assume  $E_{\text{in},j}\widehat{\ell}_j = 0$  for all  $j \in [m]$ , since otherwise we are in Case (i). Now, again ignoring all but the  $(j, z)$  term in the CE loss for which (37) holds for some  $v \notin \mathcal{S}_j$ , we find

$$\text{CE}(\widehat{\mathbf{L}}_B) - \mathcal{H} \geq \hat{\pi}_j \hat{p}_{j,z} \log \left( p_{j,z} \left( \sum_{z' \in \mathcal{S}_j} e^{-(\widehat{\mathbf{L}}_B[j,z] - \widehat{\mathbf{L}}_B[j,z'])} + e^{(\widehat{\mathbf{L}}_B[j,v] - \widehat{\mathbf{L}}_B[j,z])} \right) \right).$$

On the one hand, using  $E_{\text{in},j}\widehat{\ell}_j = 0$ , we have

$$\sum_{z' \in \mathcal{S}_j} e^{-(\widehat{\mathbf{L}}_B[j,z] - \widehat{\mathbf{L}}_B[j,z'])} = \sum_{z' \in \mathcal{S}_j} e^{-(\mathbf{L}^{\text{in}}[j,z] - \mathbf{L}^{\text{in}}[j,z'])} = \sum_{z' \in \mathcal{S}_j} \frac{\hat{p}_{j,z'}}{\hat{p}_{j,z}} = \frac{1}{\hat{p}_{j,z}}.$$

On the other hand, by (37):

$$e^{\widehat{\mathbf{L}}_B[j,v] - \widehat{\mathbf{L}}_B[j,z]} \geq e^{-R'(1-\delta)} e^{\mathbf{L}^{\text{in}}[j,v] - \mathbf{L}^{\text{in}}[j,z]} \geq e^{-R'(1-\delta)} \exp(-2\|\mathbf{L}^{\text{in}}\|_2),$$

Putting the above together yield:

$$\text{CE}(\widehat{\mathbf{L}}_B) - \mathcal{H} \geq \hat{\pi}_j \hat{p}_{j,z} \log \left( 1 + \frac{e^{-R'(1-\delta)} \exp(-2\|\mathbf{L}^{\text{in}}\|_2)}{V} \right) \geq \frac{e^{-R'(1-\delta)}}{mV^2 \exp(2\|\mathbf{L}^{\text{in}}\|_2) + mV}.$$

where the second inequality uses  $\log(1+x) \geq \frac{x}{1+x}, x > 0$ .

Compare this with (33). As  $B \rightarrow \infty$ , and noting that  $R, R'$  grow at the same rate, it holds  $\frac{e^{-R'(1-\delta)}}{mV^2 \exp(2\|\mathbf{L}^{\text{in}}\|_2) + mV} > V \exp(2\|\mathbf{L}^{\text{in}}\|_2) e^{-R}$ . Thus,  $\text{CE}(\widehat{\mathbf{L}}_B) > \text{CE}(\mathbf{L}_R)$ , a contradiction.

In either case, we arrive at a contradiction, which completes the proof.

#### B.4 Proof of Proposition 1

We will prove that for  $j, j' \in [m]$  with same support set  $\mathcal{S}_j = \mathcal{S}_{j'}$  the  $j$ -th and  $j'$ -th columns of  $\mathbf{L}^{\text{mm}} = [\ell_1, \dots, \ell_m]$  are same. For the sake of contradiction, suppose that this is not the cases, i.e.  $\ell_j \neq \ell_{j'}$ . Consider the following three candidate solutions of **NTP-SVM\***:


$$\begin{aligned} \mathbf{L}^{\text{mm}} &= [\ell_1, \dots, \ell_j, \dots, \ell_{j'}, \dots, \ell_m] \\ \mathbf{L}_{\text{flip}}^{\text{mm}} &= [\ell_1, \dots, \ell_{j'}, \dots, \ell_j, \dots, \ell_m] \\ \mathbf{L}_{\text{avg}}^{\text{mm}} &= [\ell_1, \dots, \frac{\ell_j + \ell_{j'}}{2}, \dots, \frac{\ell_j + \ell_{j'}}{2}, \dots, \ell_m]. \end{aligned}$$

By assumption  $\mathbf{L}^{\text{mm}}$  is optimal in **NTP-SVM\***. We now show that the other two matrices are also optimal.

Firstly, note that  $\mathbf{L}_{\text{flip}}^{\text{mm}}$  is feasible in **NTP-SVM\*** because: (i) the affine constraints only depend on the support sets, which are same for  $j, j'$ s, (ii)  $\text{rank}(\mathbf{L}_{\text{flip}}^{\text{mm}}) = \text{rank}(\mathbf{L}^{\text{mm}}) \leq d$  since permuting columns is rank-preserving operation. Moreover, permuting columns also preserves nuclear-norm of a matrix. To see this note that  $\mathbf{L}_{\text{flip}}^{\text{mm}} = \mathbf{L}^{\text{mm}} \mathbf{P} = \mathbf{U} \mathbf{\Sigma} (\mathbf{P}^\top \mathbf{V})^\top$  for some permutation matrix  $\mathbf{P} \in \mathbb{R}^{m \times m}$ . Since  $\mathbf{V}$  is partial orthonormal, the same is true for  $\mathbf{P}^\top \mathbf{V}$ . Thus,  $\|\mathbf{L}_{\text{flip}}^{\text{mm}}\|_* = \text{tr}(\mathbf{\Sigma}) = \|\mathbf{L}^{\text{mm}}\|_*$ , which proves optimality of  $\text{rank}(\mathbf{L}_{\text{flip}}^{\text{mm}})$ .

Similarly, we can argue that  $\mathbf{L}_{\text{avg}}^{\text{mm}}$  is feasible in **NTP-SVM\*** because: (i) the affine constraints only depend on the support sets, which are same for  $j, j'$ s, (ii)  $\text{rank}(\mathbf{L}_{\text{avg}}^{\text{mm}}) = \text{rank}(\mathbf{L}^{\text{mm}}) \leq d$  since adding columns and multiplying them by scalars are rank-preserving operations. But now, from convexity of the objective function, and noting that  $\mathbf{L}_{\text{avg}}^{\text{mm}} = \frac{1}{2} (\mathbf{L}^{\text{mm}} + \mathbf{L}_{\text{flip}}^{\text{mm}})$  we have

$$\|\mathbf{L}_{\text{avg}}^{\text{mm}}\|_* \leq \frac{1}{2} \|\mathbf{L}^{\text{mm}}\|_* + \frac{1}{2} \|\mathbf{L}_{\text{flip}}^{\text{mm}}\|_* = \|\mathbf{L}^{\text{mm}}\|_*.$$



**Figure 9:** Geometry properties illustration for Proposition 2. Shown two values of  $V$  for varying values of  $k \in [V - 1]$ . (Left) angles between context and word embeddings. (Right) normalized norm ratio of word to context embeddings.

Thus,  $L_{\text{avg}}^{\text{mm}}$  is also optimal.

Now, from Lemma 1, we know that for some partial orthonormal matrix  $\mathbf{R}$ , in the limit of vanishing regularization:  $\mathbf{H}_\lambda \propto \mathbf{R}^\top \Sigma^{\frac{1}{2}} \mathbf{V}^\top$  where  $\mathbf{V}, \Sigma$  are SVD factors of  $L_{\text{avg}}^{\text{mm}}$ . Since the  $j, j'$  columns of  $L_{\text{avg}}^{\text{mm}}$  are the same by construction, the same is true for the  $j, j'$  columns of  $\mathbf{V}^\top$ . This proves the desired.

## C Additional discussions on the solution of $\text{NTP-SVM}_*$

In this section, we complement the results of Sec. 4.4 on the properties of the  $\text{NTP-SVM}_*$  solution.

### C.1 Special case: support sets of equal sizes

Recall that in Prop. 2, we introduced a special case of the sparsity pattern of the support sets for which we could find the optimal solution of  $\text{NTP-SVM}_*$  in closed-form. The proposition below provides a more detailed characterization of the implicit geometry of the embeddings.

**Proposition 5.** *Assume the setting of Proposition 2. The geometry of context and word embeddings is described by the following relations:*

$$\forall v \neq v' \in [V] : \cos(\mathbf{w}_v, \mathbf{w}_{v'}) = \frac{-1}{V-1} \quad \text{and} \quad \|\mathbf{w}_v\| = \|\mathbf{w}_{v'}\| \quad (40a)$$

$$\forall j \neq j' \in [m] : \cos(\mathbf{h}_j, \mathbf{h}_{j'}) = \frac{|\mathcal{S}_j \cap \mathcal{S}_{j'}| - \frac{k^2}{V}}{k - \frac{k^2}{V}} \quad \text{and} \quad \|\mathbf{h}_j\| = \|\mathbf{h}_{j'}\| \quad (40b)$$


$$\forall j \in [m], v \in [V] : \cos(\mathbf{w}_v, \mathbf{h}_j) = \begin{cases} \sqrt{\frac{V-1}{k(V-k)}} & v \in \mathcal{S}_j \\ \frac{-1}{\sqrt{k(V-k)(V-1)}} & v \notin \mathcal{S}_j \end{cases} \quad (40c)$$

$$\forall j \in [m], v \in [V] : \frac{\|\mathbf{w}_v\|^2}{\|\mathbf{h}_j\|^2} = \frac{(V-1) \binom{V-2}{k-1}}{k(V-k)}. \quad (40d)$$

Fig. 9 visualizes these properties. For a brief discussion on this proposition, see Sec. 4.4.1. We defer the proof of these two propositions to App. C.4.

### C.2 A simple candidate solution

Recall that in Sec. 4.4, we introduced  $\tilde{\mathbf{S}} = (\mathbb{I}_V - \frac{1}{V} \mathbf{1}_V \mathbf{1}_V^\top) \mathbf{S}$  as a candidate solution for the  $\text{NTP-SVM}_*$ . Further recall that Prop. 3 provides a dual certificate condition that specifies whether  $\tilde{\mathbf{S}}$  is optimal. Our empirical evaluations show that  $\tilde{\mathbf{S}}$  successfully solves  $\text{NTP-SVM}_*$  in numerous realizations of  $\mathbf{S}$  we have examined. However, this is not the case in all realizations of the support set structure.



**Figure 11:** Same as Fig. 10 for the support set configuration in the Simplified TinyStories of Sec. 5.

In Fig. 10, we present an example of two realizations of  $S$ , where the condition in Prop. 3 is met in one instance but not in the other. The second and third column of the figure display the heatmap of  $\tilde{S}$  and the optimal solution of  $\text{NTP-SVM}_*$  as determined by CVXPY (Grant & Boyd, 2014) for the support sets  $S$  shown in the first column. Notably, although  $\tilde{S}$  is not optimal in the second row, it provides a close and simple approximation of the optimal  $L^{\text{mm}}$  which almost captures the overall structure of the optimal solution found by CVXPY. Additionally, in Fig. 11, we show similar heatmaps of the optimal logits for the setup of the Simplified TinyStories dataset in Sec. 5. Once again, while  $\tilde{S}$  is not necessarily optimal, it serves as a close proxy for the optimal solution. Note that in the case of one-hot classification under STEP imbalanced data,  $\tilde{S}$  is previously shown to be the unique minimizer of  $\text{NTP-SVM}_*$  (Thrampoulidis et al., 2022).

### C.3 Proof of Proposition 3

To prove the proposition, we first derive the dual of  $\text{NTP-SVM}_*$ .

We begin with a more convenient, but equivalent, formulation of  $\text{NTP-SVM}_*$ . One way to do this is by recalling that  $\text{NTP-SVM}_*$  is a convex relaxation to

$$\min_{W, H} \frac{1}{2} \|W\|_F^2 + \frac{1}{2} \|H\|_F^2, \quad \text{s.t.} \quad \forall j \in [m], z_j \neq z \in \mathcal{S}_j : (\mathbf{e}_{z_j} - \mathbf{e}_z)^\top W \mathbf{h}_j = 0$$

$$\forall j \in [m], v \notin \mathcal{S}_j : (\mathbf{e}_{z_j} - \mathbf{e}_v)^\top W \mathbf{h}_j \geq 1.$$

where, for each sample  $j$ , we choose an anchor label  $z_j \in \mathcal{S}_j$ . We can now relax this in terms of the matrix  $X = \begin{bmatrix} W \\ H^\top \end{bmatrix} \begin{bmatrix} W^\top & H \end{bmatrix}$ . (Note that  $L$  essentially corresponds to the off-diagonal blocks of  $X$ ). This gives us the following reformulation of  $\text{NTP-SVM}_*$ . (Formally, it is easy to show that (41) is the dual of the dual of  $\text{NTP-SVM}_*$ . This result about nuclear minimization is classic (e.g., Recht et al. (2010)) and we omit the details.):

$$\min_{X \succeq 0} \frac{1}{2} \text{tr}(X) \quad \text{s.t.} \quad \forall j \in [m], z_j \neq z \in \mathcal{S}_j : X[z_j, V+j] - X[z', V+j] = 0 \quad (41)$$

$$\forall j \in [m], v \notin \mathcal{S}_j : X[z_j, V+j] - X[v, V+j] \geq 1.$$

Define  $A_{V \times m}$  such that  $A[z, j] = -\lambda_{j,z}$ ,  $z \neq z_j$  and  $A[z_j, j] = \sum_{z \neq z_j} \lambda_{j,z}$ . ( $\mathbf{1}_V^\top A = 0$ ).

The dual problem is

$$\begin{aligned} & \max_A \min_{X \succeq 0} \sum_{j \in [m]} \sum_{v \notin \mathcal{S}_j} \lambda_{j,v} + \frac{1}{2} \operatorname{tr} \left( \begin{bmatrix} \mathbb{I} & -A \\ -A^\top & \mathbb{I} \end{bmatrix} X \right) \quad \text{s.t.} \quad \mathbf{1}_V^\top A = 0 \quad A[v, j] \leq 0, v \notin \mathcal{S}_j \\ & = \max_A \sum_{j \in [m]} \sum_{v \notin \mathcal{S}_j} \lambda_{j,v} \quad \text{s.t.} \quad \mathbf{1}_V^\top A = 0 \quad A[v, j] \leq 0, v \notin \mathcal{S}_j, \quad \|A\|_2 \leq 1 \\ & = \max_A \operatorname{tr}(A^\top \tilde{S}) \quad \text{s.t.} \quad \mathbf{1}_V^\top A = 0 \quad A[v, j] \leq 0, v \notin \mathcal{S}_j, \quad \|A\|_2 \leq 1 \end{aligned}$$

where here  $\tilde{S}$  is defined in Prop. 2. The last line above holds since,

$$\begin{aligned} \sum_{j \in [m]} \sum_{v \notin \mathcal{S}_j} \lambda_{j,v} &= \sum_{j \in [m]} \sum_{v \notin \mathcal{S}_j} \lambda_{j,v} + \sum_{j \in [m]} \sum_{\substack{v \in \mathcal{S}_j \\ v \neq z_j}} \lambda_{j,v} - \sum_{j \in [m]} \sum_{\substack{v \in \mathcal{S}_j \\ v \neq z_j}} \lambda_{j,v} \\ &= \sum_{j \in [m]} \sum_{v \neq z_j} \lambda_{j,v} - \sum_{j \in [m]} \sum_{\substack{v \in \mathcal{S}_j \\ v \neq z_j}} \lambda_{j,v} \\ &= \sum_{j \in [m]} \left( \frac{1 - |\mathcal{S}_j|}{V} + \frac{|\mathcal{S}_j|}{V} \right) \sum_{v \neq z_j} \lambda_{j,v} - \sum_{j \in [m]} \sum_{\substack{v \in \mathcal{S}_j \\ v \neq z_j}} \lambda_{j,v} \\ &= \sum_{j \in [m]} \frac{1 - |\mathcal{S}_j|}{V} A[z_j, j] - \sum_{j \in [m]} \frac{|\mathcal{S}_j|}{V} \sum_{v \neq z_j} A[v, j] + \sum_{j \in [m]} \sum_{\substack{v \in \mathcal{S}_j \\ v \neq z_j}} A[v, j] \\ &= \sum_{j \in [m]} \tilde{S}[z_j, j] A[z_j, j] + \sum_{j \in [m]} \sum_{v \notin \mathcal{S}_j} \tilde{S}[v, j] A[v, j] + \sum_{j \in [m]} \sum_{\substack{v \in \mathcal{S}_j \\ v \neq z_j}} \tilde{S}[v, j] A[v, j] \\ &= \sum_{j \in [m]} \sum_{v \in [V]} \tilde{S}[v, j] A[v, j]. \end{aligned}$$

We remark the following about the dual problem derived above. First, because  $\tilde{S}$  is feasible in **NTP-SVM<sub>\*</sub>**, and constraints are linear, Slater's conditions hold, thus, we have strong duality. Second, by complementary slackness, if the nonnegativity conditions  $A[v, j] \leq 0, v \notin \mathcal{S}_j$  are strict, then, at optimality,  $L$  satisfies the inequality constraints (10) with equality.

Now, regarding solving the dual, observe that if we remove the constraint  $A[v, j] \leq 0, v \notin \mathcal{S}_j$ , then the solution to the dual is simply  $A = \mathbf{U}\mathbf{V}^\top$  where  $\mathbf{U}\mathbf{\Sigma}\mathbf{V}^\top$  is the SVD of  $\tilde{S}$ . In fact, it is not hard to show that this is the unique solution of the relaxed dual; see for example (Thrampoulidis et al., 2022, Lemma C.1) for an analogous result. Therefore, if  $\mathbf{U}\mathbf{V}^\top$  satisfies the condition of the proposition, then it is clearly the unique optimal solution to the dual problem.

From this and complementary slackness discussed above, it must be that any minimizer of **NTP-SVM<sub>\*</sub>** must satisfy the inequality constraints with equalities. Formally, the optimal solution  $\hat{L}$  satisfies simultaneously:

$$\hat{L}^\top \mathbf{1}_V = 0_m \quad \text{and} \quad \forall j \in [m] : \mathbf{E}_{\text{in},j} \hat{\ell}_j = \mathbf{1}_{\mathcal{S}_j}, \quad \mathbf{E}_{\text{out},j} \hat{\ell}_j = 0_{V-1-S_j}.$$

Note that this is a system of  $m + \sum_{j \in [m]} (S_j + V - 1 - S_j) = Vm$  linear equations. It can also be easily verified that these equations are linearly independent. Thus, their unique solution is  $\tilde{S}$ . This completes the proof of the proposition.

#### C.4 Proof of Propositions 2 and 5

At the heart of proving both Propositions 2 and 5 is the following result, which determines the SVD factors  $\mathbf{U}$ ,  $\mathbf{\Sigma}$ , and  $\mathbf{V}$  of  $\tilde{S} = (\mathbb{I}_V - \frac{1}{V} \mathbf{1}\mathbf{1}^\top) \mathbf{S}$ . This shows immediately that word

embeddings form an ETF. It also forms the basis for computing the geometry of context embeddings. We show this after proving the lemma.

**Lemma 8.** Fix any  $k \in [V]$  and suppose  $\mathcal{S}$  contains all  $\binom{V}{k}$  support sets of size  $k$ , i.e.  $m = \binom{V}{k}$ . Let  $\mathbf{P} \in \mathbb{R}^{V \times (V-1)}$  be orthonormal basis of the subspace orthogonal to  $\mathbf{1}_V$ . Then, the  $\mathbf{U}, \mathbf{\Sigma}$  factors of the SVD of  $\tilde{\mathcal{S}}$  are:  $\mathbf{U} = \mathbf{P}$  and  $\mathbf{\Sigma} = \sqrt{\binom{V-2}{k-1}} \mathbb{I}_{V-1}$ . Thus,

$$\mathbf{U}\mathbf{\Sigma}\mathbf{U}^\top \propto \left( \mathbb{I}_V - \frac{1}{V} \mathbf{1}\mathbf{1}^\top \right).$$

Moreover, the matrix  $\mathbf{V} \in \mathbb{R}^{m \times (V-1)}$  of right-singular vectors is such that its rows  $\mathbf{v}_j, j \in [m]$  can be expressed as follows with respect to the rows  $\mathbf{u}_z, z \in [V]$  of  $\mathbf{U}$ :

$$\mathbf{v}_j = \frac{1}{\sqrt{\binom{V-2}{k-1}}} \sum_{z \in \mathcal{S}_j} \mathbf{u}_z \quad j \in [m].$$

*Proof.* We will compute explicitly  $\tilde{\mathcal{S}}\tilde{\mathcal{S}}^\top$ . Start with computing  $\mathcal{S}\mathcal{S}^\top$ . For diagonal elements, the dot product is between a row of  $\mathcal{S}$  and itself, counts the number of 1s in that row. Since each element is included in  $\binom{V-1}{k-1}$  support sets (choosing the remaining  $k-1$  elements from the other  $V-1$  elements), each diagonal entry of  $\mathcal{S}\mathcal{S}^\top$  is  $\binom{V-1}{k-1}$ . For off-diagonal elements, the dot product counts the number of support sets in which both elements of the corresponding rows appear. This requires choosing the remaining  $k-2$  elements from  $V-2$  elements, giving  $\binom{V-2}{k-2}$  for each off-diagonal entry of  $\mathcal{S}\mathcal{S}^\top$ . Overall, after algebraic simplification, we find that

$$\mathcal{S}\mathcal{S}^\top = \binom{V-2}{k-1} \left( \mathbb{I}_V + \frac{k-1}{V-k} \mathbf{1}\mathbf{1}^\top \right).$$

We may now use the closed form of  $\tilde{\mathcal{S}}$  and few more algebra simplifications to find that

$$\tilde{\mathcal{S}}\tilde{\mathcal{S}}^\top = \binom{V-2}{k-1} \left( \mathbb{I}_V - \frac{1}{V} \mathbf{1}\mathbf{1}^\top \right).$$

Let  $\mathbf{P} \in \mathbb{R}^{V \times (V-1)}$  denote an orthonormal basis of the subspace orthogonal to  $\mathbf{1}_V$ . Since  $\tilde{\mathcal{S}}\tilde{\mathcal{S}}^\top = \binom{V-2}{k-1} \mathbf{P}\mathbf{P}^\top$  and  $\tilde{\mathcal{S}}\tilde{\mathcal{S}}^\top = \mathbf{U}\mathbf{\Sigma}^2\mathbf{U}^\top$ , we have shown that  $\mathbf{U} = \mathbf{P}$  and  $\mathbf{\Sigma} = \sqrt{\binom{V-2}{k-1}} \mathbb{I}_V$ .  $\square$

Using  $\mathbf{1}^\top \mathcal{S} = k\mathbf{1}_m^\top$ , we can compute

$$\tilde{\mathcal{S}}^\top \tilde{\mathcal{S}} = \mathcal{S}^\top \mathcal{S} - \frac{1}{V} (\mathcal{S}^\top \mathbf{1}_V) (\mathcal{S}^\top \mathbf{1}_V)^\top = \mathcal{S}^\top \mathcal{S} - \frac{k^2}{V} \mathbf{1}_m \mathbf{1}_m^\top.$$

From this and the fact that  $\text{diag}(\mathcal{S}^\top \mathcal{S}) = k\mathbb{I}_m$ , and recalling that  $\tilde{\mathcal{S}}^\top \tilde{\mathcal{S}} = \mathbf{V}\mathbf{\Sigma}^2\mathbf{V}^\top$ , we find that all context embeddings are equinorm with squared norm (i.e. diagonal entries of  $\mathbf{V}\mathbf{\Sigma}\mathbf{V}^\top$ ) equal to

$$\|\mathbf{h}_j\|^2 = \frac{k - \frac{k^2}{V}}{\sqrt{\binom{V-2}{k-1}}}, \quad \forall j \in [m].$$

Recall also that the diagonal entries of  $\mathbf{U}\mathbf{\Sigma}\mathbf{U}^\top$  are

$$\|\mathbf{w}_v\|^2 = \left( 1 - \frac{1}{V} \right) \sqrt{\binom{V-2}{k-1}}, \quad \forall v \in [V].$$

From the off-diagonal entries of  $\tilde{\mathcal{S}}^\top \tilde{\mathcal{S}}$ , we can infer the angles between different context embeddings. Let  $j \neq j' \in [m]$ , then the  $(j, j')$  of  $\tilde{\mathcal{S}}^\top \tilde{\mathcal{S}}$  is  $|\mathcal{S}_j \cap \mathcal{S}_{j'}| - \frac{k^2}{V}$ . Thus, the  $(j, j')$  off-diagonal entries of  $\mathbf{V}\mathbf{\Sigma}\mathbf{V}^\top$  are

$$\mathbf{h}_j^\top \mathbf{h}_{j'} = \frac{|\mathcal{S}_j \cap \mathcal{S}_{j'}| - \frac{k^2}{V}}{\sqrt{\binom{V-2}{k-1}}} \implies \cos(\mathbf{h}_j, \mathbf{h}_{j'}) = \frac{|\mathcal{S}_j \cap \mathcal{S}_{j'}| - \frac{k^2}{V}}{k - \frac{k^2}{V}} \quad j \neq j' \in [m].$$



Recalling that  $\mathbf{W}\mathbf{H}^\top = \mathbf{L}^{\text{mm}}$ , we may now compute the angle between word and context embeddings as follows:

$$\cos(\mathbf{w}_v, \mathbf{h}_j) = \begin{cases} \sqrt{\frac{V-1}{k(V-k)}} & v \in \mathcal{S}_j \\ \frac{-1}{\sqrt{k(V-k)(V-1)}} & v \notin \mathcal{S}_j \end{cases}. \quad (42)$$

This completes the proof of Proposition 5 and the proof of statements (ii) and (iii) of Proposition 2.

It remains to prove statement (i) of Proposition 2, i.e. proving that  $\tilde{\mathbf{S}} = (\mathbb{I}_V - \frac{1}{V}\mathbf{1}\mathbf{1}^\top)\mathbf{S}$  is the unique minimizer of  $\text{NTP-SVM}_*$ .

To prove this, we appeal to Proposition 3. First, it is straightforward checking that  $\mathbf{L}^{\text{mm}}$  satisfies the SVM constraints (in fact, with equality). Thus, from Proposition 3, it suffices that the matrix  $\mathbf{A} := \mathbf{U}\mathbf{V}^\top$  satisfies  $\mathbf{A}[v, j] < 0$  for all  $v \notin \mathcal{S}_j$  and all  $j \in [m]$ . But, in our case  $\mathbf{A} = \mathbf{U}\mathbf{V}^\top = \binom{V-2}{k-1}^{-\frac{1}{2}}\tilde{\mathbf{S}}$  because  $\mathbf{\Sigma} = \mathbb{I}_{V-1}$ . Thus, the desired condition holds for  $\mathbf{A}$  since  $\tilde{\mathbf{S}}[v, j] = -k/V < 0$  for all  $j \notin \mathcal{S}_j$  and  $j \in [m]$ .

## D Numerical experiments

### D.1 Additional details and results: Sec. 5

#### D.1.1 Datasets

We use a total of three datasets. We first employ two smaller-scale datasets, which are well-suited for verifying our theoretical solutions. Then, we use one larger-scale dataset designed to investigate the geometric properties of text in conditions that approximate real-world text scenarios.

**Synthetic.** We generate each context  $\bar{x}_j$  as follows: We select  $T - 1 = 5$  words (tokens), and manually come up with the  $T = 6$ -th tokens that are consistent with the given context. To model the probabilities  $\hat{\pi}_j$  and  $\hat{p}_j$ , we sample each support independently to a number of repeats to emulate the behavior of repeated context in natural language. For instance, the context  $\bar{x}_j = \text{“Lily wants to try the”}$  is followed by  $\mathcal{S} = \{\text{soup, game, movie}\}$  with respective probabilities  $[0.5, 0.25, 0.25]$ . The dataset consists of  $n = 116$  samples, containing  $m = 16$  distinct contexts, and a vocabulary size of  $V = 30$ . Each context has a fixed support set length of  $|\mathcal{S}_j| = 3$ . In Fig. 12, we show the  $m = 16$  unique contexts and the soft labels on their next token. In this dataset, the tokenization is done at the word level and the empirical entropy lower-bound is  $\mathcal{H} = 1.6597$ .

**Simplified TinyStories.** Advancing towards a more realistic but still controlled dataset, we use contexts and support sets from the TinyStories corpus. We derive contexts  $\bar{x}_j$  and support sets  $\mathcal{S}_j$  from the TinyStories dataset and we create the training data from the most frequent word-level contexts with length 5, such as  $\bar{x}_j = [\text{“once”, “upon”, “a”, “time”, “,”}]$ . For the support sets  $\mathcal{S}_j$ , we record all next tokens of  $\bar{x}_j$  in the original dataset. Then, we replace the words in the contexts with their synonyms to generate new contexts with identical support sets. This allows us to have multiple contexts sharing common support sets, while controlling the vocabulary size, which is set to  $V = 104$ . The final dataset consists of  $n \approx 3050$  samples with  $m \approx 400$  unique contexts. The frequencies  $\hat{p}_j$  are determined by independently sampling each next token from  $\mathcal{S}_j$  several times. Here,  $\mathcal{H} = 1.0166$ .

**TinyStories.** To experiment with more a standard dataset, we use 100 stories sampled from TinyStories. Here, we do not make any sample selections for contexts and support sets, unlike the previously mentioned datasets. We use a tokenizer with vocabulary size  $V = 128$ , and fixed context length  $T - 1 = 6$ . We choose a small context window and vocabulary size to make tracking the distinct contexts and their support sets computationally manageable. In this setup, the number of distinct contexts  $m \sim 10^5$  and  $\mathcal{H} = 0.3112$ .

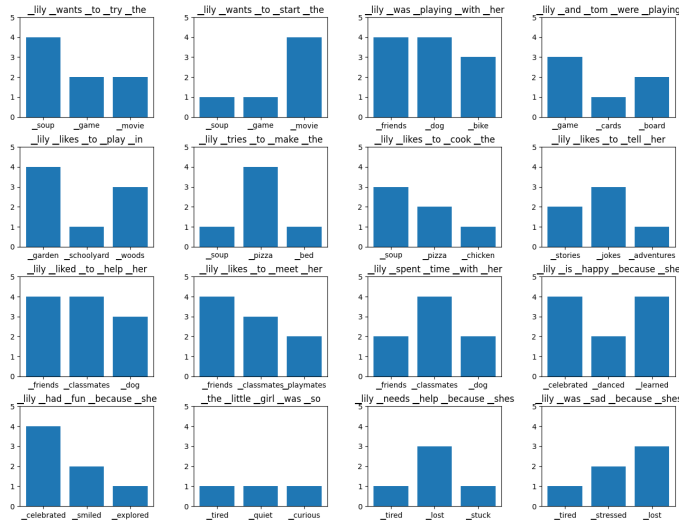


Figure 12: The contexts, support sets, and soft labels in the Synthetic dataset.

### D.1.2 Training details

**Synthetic and TinyStories experimtns.** We train the models long enough to ensure that the loss approached the empirical entropy lower-bound  $\mathcal{H}$  within an order of  $10^{-4}$ . In all experiments, we use AdamW optimizer with a weight decay  $\lambda = 10^{-6}$ . The learning rate is initially set to  $10^{-4}$ , with a step-decay schedule.

**TinyStories experimtns.** We train TF with AdamW optimizer until it reaches the empirical entropy lower-bound within an order  $10^{-3}$ . We use warm-up over the first 5 epochs to increase the learning rate to  $5 \times 10^{-4}$  and use a cosine learning rate scheduler to decay its values over the course of the training. We set the weight decay to  $10^{-5}$ .

### D.1.3 Additional results

Figs. 13-16 complement the discussions and experimental results in Sec. 5. Below we only discuss the missing details of the TinyStories experiments.

**Visualization details of TinyStories experiments.** For visualizing the heatmap of embeddings similarity in Fig. 3, we choose 10 unique support sets in the dataset with  $|\mathcal{S}_j| > 2$ . For each of the chosen support sets  $\mathcal{S}$ , we choose 10 distinct contexts  $j$  such that  $\tilde{\mathcal{S}}_j = \mathcal{S}$  (a total of 100 samples), and illustrate the Gram matrix of the centered support matrix  $\tilde{\mathbf{S}}$  and context embeddings  $\mathbf{H}$  for this subset of the contexts. Given that the support sets are sparse, we choose to use contexts with the largest support set size which would result in a more pronounced pattern on the correlation matrices. We skip visualizing the word embeddings  $\mathbf{W}$  in this experiment, as the heatmaps did not display any noticeable visual structures.

For the metrics displayed in Fig. 7, due to the larger number of context embeddings, we only compute the metrics such as  $\|\mathcal{P}_{\mathcal{F}}(\mathbf{L}_k) - \mathbf{L}^{\text{in}}\|$ , norm growth and correlation with proxies using a sample of 1000 token/context pairs chosen randomly from the training set. We observe qualitatively similar behavior compared to other experiments, with the exception of the norm growth of  $\mathbf{H}$ , which we conjecture to be a consequence of layer normalization in the final embedding layer.

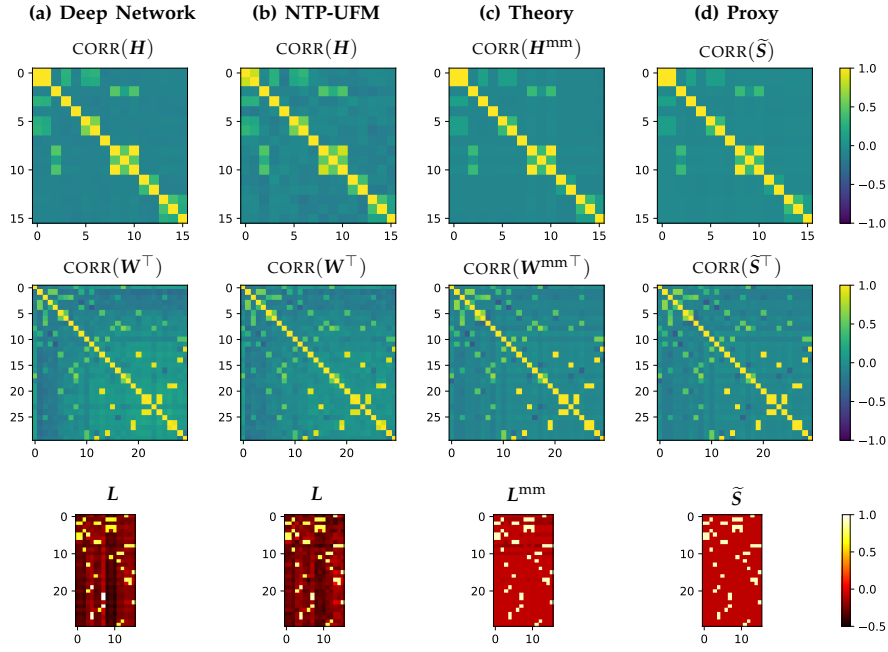


Figure 13: Same as Fig. 1, this time for the Synthetic dataset.

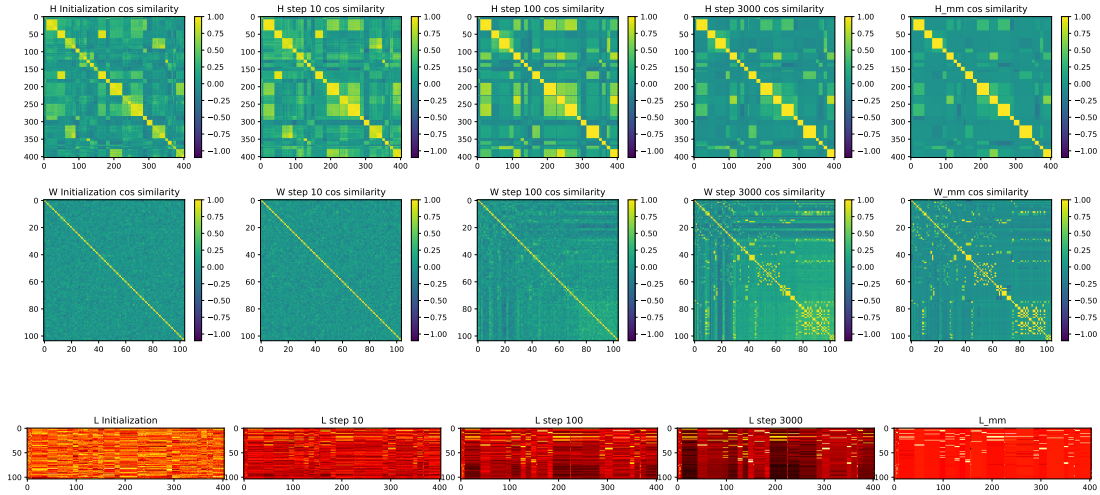


Figure 14: Evolution of  $H_k$ ,  $W_k$  and  $L_k$  to NTP-SVM<sub>\*</sub> solution  $H^{mm}$ ,  $W^{mm}$  and  $L^{mm}$  throughout training. The embeddings are trained by the transformer on the Simplified TinyStories dataset. See Sec. 5.

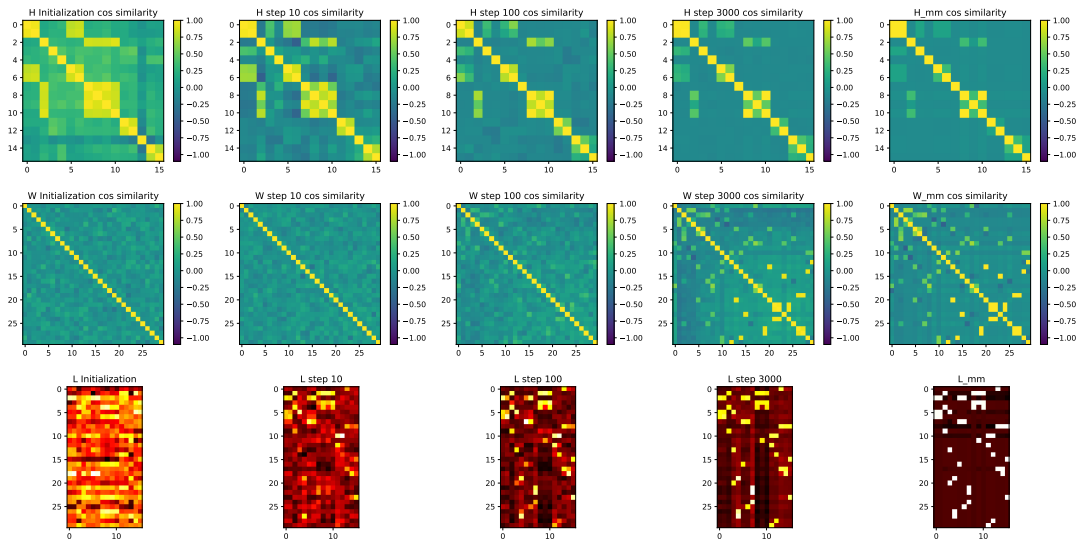
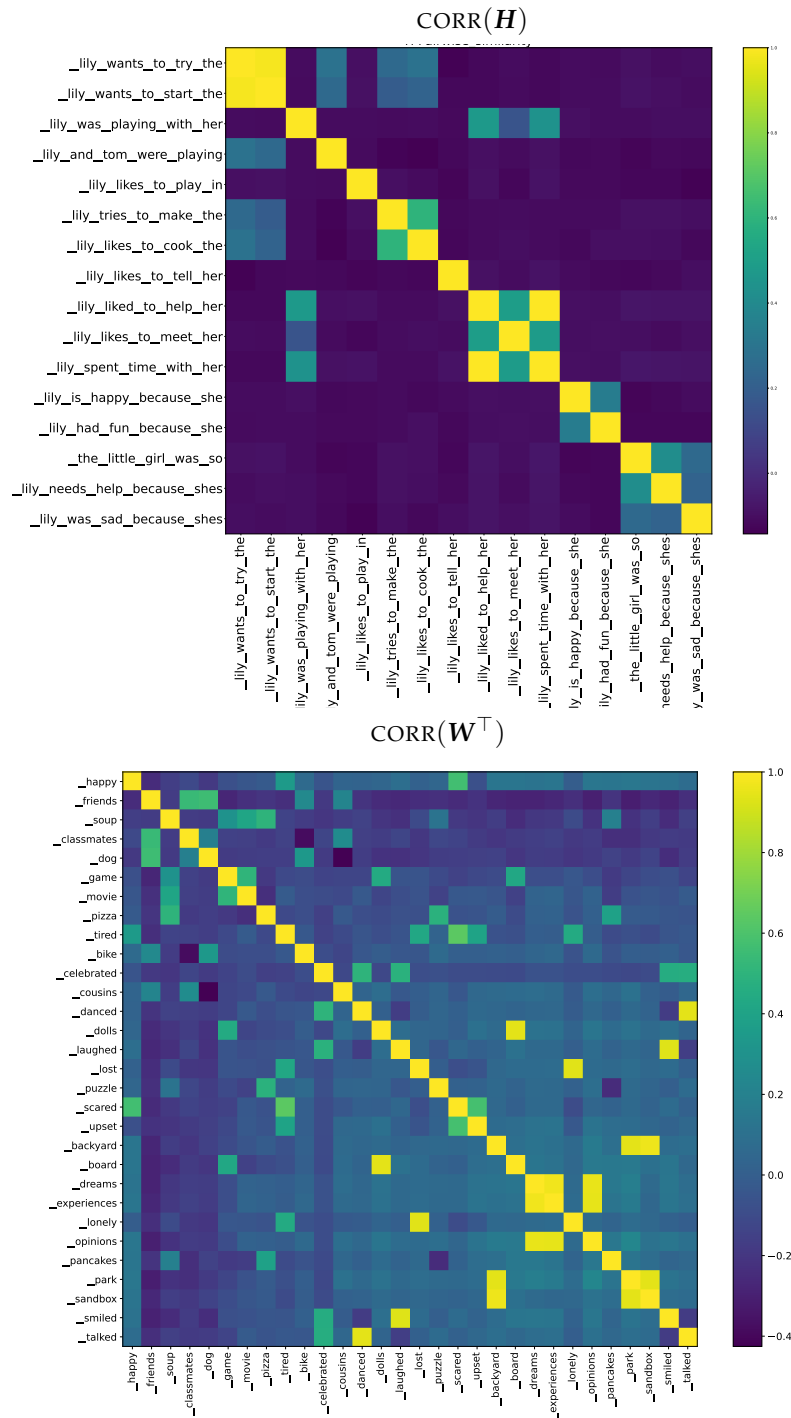


Figure 15: Same as Fig. 14, this time for the Synthetic dataset. See Sec. 5.



**Figure 16:** Geometry of context and word embeddings for Synthetic dataset. A lighter color indicates higher similarity in the embedding space.

## D.2 Other architectures

The results in Thms. 1 and 2 hold provided the model is expressive enough to generate (approximately) unconstrained embeddings and the loss can be minimized close to the lower-bound. Depending on the expressiveness of a specific network design, the size of the network required for achieving the entropy lower-bound can vary significantly.

To explore this, we hereby repeat our experiments on Simplified TinyStories by replacing the TF model with a multi-layer perceptron (MLP). We use an MLP consisting of 20 layers organized into four blocks, each containing five layers, with hidden dimension sizes 1024, 512, 256, and 128. This leads to a network with around 8 times more parameters than the TF model. The final embeddings geometry is depicted in Fig. 17. We observe that the geometrical patterns appearing in  $\text{CORR}(\mathbf{H})$  and  $\text{CORR}(\mathbf{W}^\top)$  at a coarse level are similar to those of the TF model. However, the MLP, even with  $\sim 10$  times larger size, struggles to recover the fine-grained patterns. In terms of loss convergence, we find the MLP converges to the empirical entropy in the order of  $10^{-2}$  but TF converges in the order of  $10^{-4}$  on the same dataset. We conjecture that an even larger MLP might be able to achieve better convergence to our theoretical prediction, but we also caution of potential optimization bottlenecks. We encourage additional experiments with other architectures such as LSTMs or state-space models as future work.




Figure 17: Embedding geometries trained by MLP. Left: Geometry comparison between  $\text{CORR}(\tilde{\mathbf{S}})$  and  $\text{CORR}(\mathbf{H})$ ; Right: Geometry comparison between  $\text{CORR}(\tilde{\mathbf{S}}^\top)$  and  $\text{CORR}(\mathbf{W}^\top)$ . See App. D.2 for details.

## D.3 Auto-regressive training

Here, we train the model auto-regressively. In the previous settings, we focused on training samples that were all of fixed length  $T - 1 = 5$ . In this section, we let the model learn the embeddings for different sequence lengths  $1 \leq T - 1 \leq 16$ . Note that the length of the context does not have an impact on our theoretical analysis other than affecting the sparsity pattern  $\mathbf{S}$  of the training set which in turn influences the implicit geometry as we have seen.

We train an 8-layer TF on 200 stories from TinyStories using character-level tokenizer, which limits the vocabulary to approximately 40 characters and promotes higher entropy in next-token distribution. We denote the loss value across different positions  $T \in \{2, \dots, 17\}$  at iteration  $k$  by  $\text{CE}(L_{T,k})$ . In Fig. 19, we display the distance of each loss component  $\text{CE}(L_{T,k})$  from its empirical entropy lower-bound  $\hat{\mathcal{H}}_T$ , i.e., the T-gram entropy of the training set. We observe generally better convergence for shorter sequence lengths  $T$ . In Fig. 18, we illustrate the context embeddings similarities. For visualization, we select 5 contexts per each sequence length that end in token “y\_” or “\_t”. The context embeddings and the proxy  $(\mathbf{P})$  show similar patterns at a coarse level. However, at finer scale, the learned context embeddings that end with the same tokens exhibit strong correlation on average, even when their support sets do not align. We defer further investigation into whether the embeddings’ dependence is due to insufficient network capacity or an optimization bottleneck to future work.

## D.4 Discussion on large vocabulary setting

Throughout our analysis, we require the embedding dimension to be larger than the vocabulary size, i.e.,  $d \geq V$ . This condition was necessary to make the non-convex problem in (8) convex. However, if there exists a low-rank optimal solution in  $\text{NTP-SVM}_*$ , i.e.,  $\text{rank}(L^{\text{mm}}) < V$ , the condition  $d \geq V$  can be relaxed in the analysis. In general, the smallest



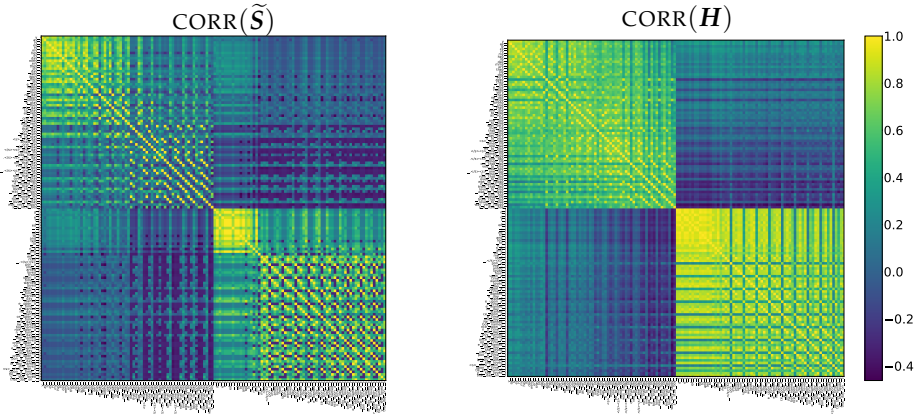


Figure 18: Geometry of context embeddings  $\text{CORR}(H)$  and the heuristic proxy  $(P)$ ,  $\text{CORR}(\tilde{S})$ , in the autoregressive experiment of App. D.3.


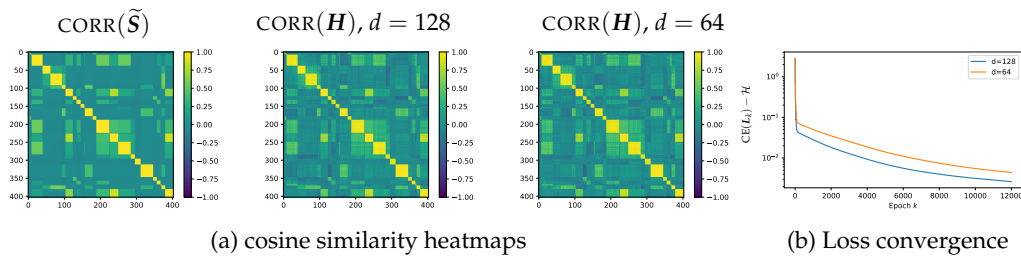


Figure 19: TF trained autoregressively on a subset of TinyStories dataset: Loss convergence to its empirical entropy lower-bound for each sequence length  $T = 2, \dots, 17$ . See App. D.3 for details.

rank among all minimizers depends on the sparsity pattern of the language. It is intriguing to investigate this dependence further.

This section examines the network’s performance when the decoder dimension  $d$  is smaller than the vocabulary size  $V$ . To ensure a fair comparison between different setups, we fix the transformer blocks’ inner embedding dimension  $d_{\text{TF}} = 64$  to maintain comparable expressivity for all networks. To vary the final layer embedding dimension  $d$ , we add an additional linear layer on top of the network to adjust the context embedding dimension.

In Fig. 20, we show the learned context embeddings for  $d = 128$  and  $d = 64$ , for a 10-layer TF trained on the Simplified TinyStories dataset with  $V = 104$ . We observe that with smaller  $d$ , the speed of convergence to the entropy lower-bound  $\mathcal{H}$  decreases. However, the learned context embeddings exhibit relatively similar geometry with moderate values of  $d < V$ . We leave more theoretical and empirical exploration of this setting to future works.



**Figure 20:** Varying embedding dimension  $d$  in the Simplified TinyStories experiment. **(a)** Left: Proxy ( $\mathbf{P}$ ) for context embeddings. Middle and Right: Geometry of context embeddings trained with TF with  $d = 128 > V$  and  $d = 64 < V$ , respectively, **(b)** Loss convergence to the empirical entropy lower bound  $\mathcal{H}$ . See App. D.4 for details.