

# Sparse Sequence-to-Sequence Models

Ben Peters<sup>†</sup> Vlad Niculae<sup>†</sup> and André F. T. Martins<sup>†‡</sup>

<sup>†</sup>Instituto de Telecomunicações, Lisbon, Portugal

<sup>‡</sup>Unbabel, Lisbon, Portugal

benzurdopeters@gmail.com, vlad@vene.ro, andre.martins@unbabel.com

## Abstract

Sequence-to-sequence models are a powerful workhorse of NLP. Most variants employ a softmax transformation in both their attention mechanism and output layer, leading to dense alignments and strictly positive output probabilities. This density is wasteful, making models less interpretable and assigning probability mass to many implausible outputs. In this paper, we propose *sparse* sequence-to-sequence models, rooted in a new family of  $\alpha$ -entmax transformations, which includes softmax and sparsemax as particular cases, and is sparse for any  $\alpha > 1$ . We provide fast algorithms to evaluate these transformations and their gradients, which scale well for large vocabulary sizes. Our models are able to produce sparse alignments and to assign nonzero probability to a short list of plausible outputs, sometimes rendering beam search exact. Experiments on morphological inflection and machine translation reveal consistent gains over dense models.

## 1 Introduction

Attention-based sequence-to-sequence (*seq2seq*) models have proven useful for a variety of NLP applications, including machine translation (Bahdanau et al., 2015; Vaswani et al., 2017), speech recognition (Chorowski et al., 2015), abstractive summarization (Chopra et al., 2016), and morphological inflection generation (Kann and Schütze, 2016), among others. In part, their strength comes from their flexibility: many tasks can be formulated as transducing a source sequence into a target sequence of possibly different length.

However, conventional *seq2seq* models are **dense**: they compute both attention weights and output probabilities with the **softmax function** (Bridle, 1990), which always returns positive values. This results in *dense attention alignments*, in which each source position is attended to at each




Figure 1: The full beam search of our best performing morphological inflection model when generating the past participle of the verb “draw”. The model gives nonzero probability to exactly three hypotheses, including the correct form (“drawn”) and the form that would be correct if “draw” were regular (“drawed”).

target position, and in *dense output probabilities*, in which each vocabulary type always has nonzero probability of being generated. This contrasts with traditional statistical machine translation systems, which are based on sparse, hard alignments, and decode by navigating through a sparse lattice of phrase hypotheses. Can we transfer such notions of sparsity to modern neural architectures? And if so, do they improve performance?

In this paper, we provide an affirmative answer to both questions by proposing **neural sparse seq2seq models** that replace the softmax transformations (both in the attention and output) by **sparse transformations**. Our innovations are rooted in the recently proposed sparsemax transformation (Martins and Astudillo, 2016) and Fenchel-Young losses (Blondel et al., 2019). Concretely, we consider a family of transformations (dubbed  $\alpha$ -entmax), parametrized by a scalar  $\alpha$ , based on the Tsallis entropies (Tsallis, 1988). This family includes softmax ( $\alpha = 1$ ) and sparsemax ( $\alpha = 2$ ) as particular cases. Crucially, entmax transforms are sparse for all  $\alpha > 1$ .

Our models are able to produce both **sparse attention**, a form of inductive bias that increases focus on relevant source words and makes alignments more interpretable, and **sparse output probabilities**, which together with auto-regressive mod-

<b>This</b>	92.9%	<b>is another</b>	view	49.8%	<b>at</b>	95.7%	<b>the tree of life .</b>
So	5.9%		<b>look</b>	27.1%	<b>on</b>	5.9%	
And	1.3%		glimpse	19.9%	,	1.3%	
Here	<0.1%		kind	2.0%			
			looking	0.9%			
			way	0.2%			
			vision	<0.1%			
			<b>gaze</b>	<0.1%			

Figure 2: Forced decoding using sparsemax attention and 1.5-entmax output for the German source sentence, “Dies ist ein weiterer Blick auf den Baum des Lebens.” Predictions with nonzero probability are shown at each time step. All other target types have probability exactly zero. When consecutive predictions consist of a single word, we combine their borders to showcase *auto-completion* potential. The selected gold targets are in boldface.

els can lead to probability distributions that are nonzero only for a finite subset of all possible strings. In certain cases, a short list of plausible outputs can be enumerated without ever exhausting the beam (Figure 1), rendering beam search **exact**. Sparse output *seq2seq* models can also be used for adaptive, sparse next word suggestion (Figure 2).

Overall, our contributions are as follows:

- We propose an **entmax sparse output layer**, together with a natural loss function. In large-vocabulary settings, sparse outputs avoid wasting probability mass on unlikely outputs, substantially improving accuracy. For tasks with little output ambiguity, entmax losses, coupled with beam search, can often produce **exact finite sets** with only one or a few sequences. To our knowledge, this is the first study of sparse output probabilities in *seq2seq* problems.
- We construct **entmax sparse attention**, improving interpretability at no cost in accuracy. We show that the entmax gradient has a simple form (Proposition 2), revealing an insightful missing link between softmax and sparsemax.
- We derive a novel exact algorithm for the case of 1.5-entmax, achieving processing speed close to softmax on the GPU, even with large vocabulary sizes. For arbitrary  $\alpha$ , we investigate a GPU-friendly approximate algorithm.<sup>1</sup>

We experiment on two tasks: one character-level with little ambiguity (**morphological inflection generation**) and another word-level, with more ambiguity (**neural machine translation**). The results

<sup>1</sup>Our standalone Pytorch entmax implementation is available at <https://github.com/deep-spin/entmax>.

show clear benefits of our approach, both in terms of accuracy and interpretability.

## 2 Background

The underlying architecture we focus on is an RNN-based *seq2seq* with global attention and input-feeding (Luong et al., 2015). We provide a brief description of this architecture, with an emphasis on the attention mapping and the loss function.

**Notation.** Scalars, vectors, and matrices are denoted respectively as  $a$ ,  $\mathbf{a}$ , and  $\mathbf{A}$ . We denote the  $d$ -probability simplex (the set of vectors representing probability distributions over  $d$  choices) by  $\Delta^d := \{\mathbf{p} \in \mathbb{R}^d : \mathbf{p} \geq 0, \|\mathbf{p}\|_1 = 1\}$ . We denote the positive part as  $[a]_+ := \max\{a, 0\}$ , and by  $[\mathbf{a}]_+$  its elementwise application to vectors. We denote the indicator vector  $\mathbf{e}_y := [0, \dots, 0, \underbrace{1}_y, 0, \dots, 0]$ .

**Encoder.** Given an input sequence of tokens  $\mathbf{x} := [x_1, \dots, x_J]$ , the encoder applies an embedding lookup followed by  $K$  layered bidirectional LSTMs (Hochreiter and Schmidhuber, 1997), resulting in encoder states  $[\mathbf{h}_1, \dots, \mathbf{h}_J]$ .

**Decoder.** The decoder generates output tokens  $y_1, \dots, y_T$ , one at a time, terminated by a stop symbol. At each time step  $t$ , it computes a probability distribution for the next generated word  $y_t$ , as follows. Given the current state  $\mathbf{s}_t$  of the decoder LSTM, an **attention mechanism** (Bahdanau et al., 2015) computes a focused, fixed-size summary of the encodings  $[\mathbf{h}_1, \dots, \mathbf{h}_J]$ , using  $\mathbf{s}_t$  as a query vector. This is done by computing token-level scores  $z_j := \mathbf{s}_t^\top \mathbf{W}^{(z)} \mathbf{h}_j$ , then taking a weighted average

$$\mathbf{c}_t := \sum_{j=1}^J \pi_j \mathbf{h}_j, \text{ where } \pi := \text{softmax}(\mathbf{z}). \quad (1)$$

The contextual output is the non-linear combination  $\mathbf{o}_t := \tanh(\mathbf{W}^{(o)}[\mathbf{s}_t; \mathbf{c}_t] + \mathbf{b}^{(o)})$ , yielding the predictive distribution of the next word

$$p(y_t = \cdot \mid \mathbf{x}, y_1, \dots, y_{t-1}) := \text{softmax}(\mathbf{V}\mathbf{o}_t + \mathbf{b}). \quad (2)$$

The output  $\mathbf{o}_t$ , together with the embedding of the predicted  $\hat{y}_t$ , feed into the decoder LSTM for the next step, in an auto-regressive manner. The model is trained to maximize the likelihood of the correct target sentences, or equivalently, to minimize

$$\mathcal{L} = \sum_{(x,y) \in \mathcal{D}} \sum_{t=1}^{|y|} \underbrace{(-\log \text{softmax}(\mathbf{V}\mathbf{o}_t))}_{L_{\text{softmax}}(y_t, \mathbf{V}\mathbf{o}_t)} y_t. \quad (3)$$

A central building block in the architecture is the transformation  $\text{softmax}: \mathbb{R}^d \rightarrow \Delta^d$ ,

$$\text{softmax}(\mathbf{z})_j := \frac{\exp(z_j)}{\sum_i \exp(z_i)}, \quad (4)$$

which maps a vector of scores  $\mathbf{z}$  into a probability distribution (*i.e.*, a vector in  $\Delta^d$ ). As seen above, the softmax mapping plays **two crucial roles** in the decoder: first, in computing normalized attention weights (Eq. 1), second, in computing the predictive probability distribution (Eq. 2). Since  $\exp \succeq 0$ , softmax never assigns a probability of zero to any word, so we may never fully rule out non-important input tokens from attention, nor unlikely words from the generation vocabulary. While this may be advantageous for dealing with uncertainty, it may be preferable to avoid dedicating model resources to irrelevant words. In the next section, we present a strategy for **differentiable sparse probability mappings**. We show that our approach can be used to learn powerful *seq2seq* models with sparse outputs and sparse attention mechanisms.

### 3 Sparse Attention and Outputs

#### 3.1 The sparsemax mapping and loss

To pave the way to a more general family of sparse attention and losses, we point out that softmax (Eq. 4) is only one of many possible mappings from  $\mathbb{R}^d$  to  $\Delta^d$ . Martins and Astudillo (2016) introduce **sparsemax**: an alternative to softmax which tends to yield **sparse probability distributions**:

$$\text{sparsemax}(\mathbf{z}) := \underset{\mathbf{p} \in \Delta^d}{\text{argmin}} \|\mathbf{p} - \mathbf{z}\|^2. \quad (5)$$

Since Eq. 5 is a projection onto  $\Delta^d$ , which tends to yield sparse solutions, the predictive distribution

$\mathbf{p}^* := \text{sparsemax}(\mathbf{z})$  is likely to assign **exactly zero** probability to low-scoring choices. They also propose a corresponding **loss function** to replace the negative log likelihood loss  $L_{\text{softmax}}$  (Eq. 3):

$$L_{\text{sparsemax}}(y, \mathbf{z}) := \frac{1}{2} (\|\mathbf{e}_y - \mathbf{z}\|^2 - \|\mathbf{p}^* - \mathbf{z}\|^2), \quad (6)$$

This loss is smooth and convex on  $\mathbf{z}$  and has a **margin**: it is zero if and only if  $z_y \geq z_{y'} + 1$  for any  $y' \neq y$  (Martins and Astudillo, 2016, Proposition 3). Training models with the sparsemax loss requires its gradient (*cf.* Appendix A.2):

$$\nabla_{\mathbf{z}} L_{\text{sparsemax}}(y, \mathbf{z}) = -\mathbf{e}_y + \mathbf{p}^*.$$

For using the sparsemax *mapping* in an attention mechanism, Martins and Astudillo (2016) show that it is differentiable almost everywhere, with

$$\frac{\partial \text{sparsemax}(\mathbf{z})}{\partial \mathbf{z}} = \text{diag}(\mathbf{s}) - \frac{1}{\|\mathbf{s}\|_1} \mathbf{s} \mathbf{s}^\top,$$

where  $s_j = 1$  if  $p_j^* > 0$ , otherwise  $s_j = 0$ .

**Entropy interpretation.** At first glance, sparsemax appears very different from softmax, and a strategy for producing other sparse probability mappings is not obvious. However, the connection becomes clear when considering the variational form of softmax (Wainwright and Jordan, 2008):

$$\text{softmax}(\mathbf{z}) = \underset{\mathbf{p} \in \Delta^d}{\text{argmax}} \mathbf{p}^\top \mathbf{z} + H^S(\mathbf{p}), \quad (7)$$

where  $H^S(\mathbf{p}) := -\sum_j p_j \log p_j$  is the well-known Gibbs-Boltzmann-Shannon entropy with base  $e$ .

Likewise, letting  $H^G(\mathbf{p}) := \frac{1}{2} \sum_j p_j (1 - p_j)$  be the **Gini entropy**, we can rearrange Eq. 5 as

$$\text{sparsemax}(\mathbf{z}) = \underset{\mathbf{p} \in \Delta^d}{\text{argmax}} \mathbf{p}^\top \mathbf{z} + H^G(\mathbf{p}), \quad (8)$$

crystallizing the connection between softmax and sparsemax: they only differ in the choice of **entropic regularizer**.

#### 3.2 A new entmax mapping and loss family

The parallel above raises a question: can we find **interesting interpolations between softmax and sparsemax**? We answer affirmatively, by considering a generalization of the Shannon and Gini entropies proposed by Tsallis (1988): a family of entropies parametrized by a scalar  $\alpha > 1$  which we call **Tsallis  $\alpha$ -entropies**:

$$H_\alpha^\top(\mathbf{p}) := \begin{cases} \frac{1}{\alpha(\alpha-1)} \sum_j (p_j - p_j^\alpha), & \alpha \neq 1, \\ H^S(\mathbf{p}), & \alpha = 1. \end{cases} \quad (9)$$




Figure 3: Illustration of entmax in the two-dimensional case  $\alpha\text{-entmax}([t, 0]_1)$ . All mappings except softmax saturate at  $t = \pm 1/\alpha - 1$ . While sparsemax is piecewise linear, mappings with  $1 < \alpha < 2$  have smooth corners.

This family is continuous, *i.e.*,  $\lim_{\alpha \rightarrow 1} H_\alpha^\top(\mathbf{p}) = H^\top(\mathbf{p})$  for any  $\mathbf{p} \in \Delta^d$  (*cf.* Appendix A.1). Moreover,  $H_2^\top \equiv H^G$ . Thus, Tsallis entropies interpolate between the Shannon and Gini entropies. Starting from the Tsallis entropies, we construct a probability mapping, which we dub **entmax**:

$$\alpha\text{-entmax}(\mathbf{z}) := \operatorname{argmax}_{\mathbf{p} \in \Delta^d} \mathbf{p}^\top \mathbf{z} + H_\alpha^\top(\mathbf{p}), \quad (10)$$

and, denoting  $\mathbf{p}^* := \alpha\text{-entmax}(\mathbf{z})$ , a loss function

$$L_\alpha(y, \mathbf{z}) := (\mathbf{p}^* - \mathbf{e}_y)^\top \mathbf{z} + H_\alpha^\top(\mathbf{p}^*) \quad (11)$$

The motivation for this loss function resides in the fact that it is a **Fenchel-Young loss** (Blondel et al., 2019), as we briefly explain in Appendix A.2. Then, 1-entmax  $\equiv$  softmax and 2-entmax  $\equiv$  sparsemax. Similarly,  $L_1$  is the negative log likelihood, and  $L_2$  is the sparsemax loss. For all  $\alpha > 1$ , entmax tends to produce **sparse probability distributions**, yielding a function family continuously interpolating between softmax and sparsemax, *cf.* Figure 3. The gradient of the entmax loss is

$$\nabla_{\mathbf{z}} L_\alpha(y, \mathbf{z}) = -\mathbf{e}_y + \mathbf{p}^*. \quad (12)$$

**Tsallis entmax losses** have useful properties including convexity, differentiability, and a hinge-like separation margin property: the loss incurred becomes zero when the score of the correct class is separated by the rest by a margin of  $1/\alpha - 1$ . When separation is achieved,  $\mathbf{p}^* = \mathbf{e}_y$  (Blondel et al., 2019). This allows entmax *seq2seq* models to be **adaptive to the degree of uncertainty** present: decoders may make fully confident predictions at “easy” time steps, while preserving **sparse uncertainty** when a few choices are possible (as exemplified in Figure 2).

**Tsallis entmax probability mappings** have not, to our knowledge, been used in attention mechanisms. They inherit the desirable sparsity of sparsemax, while exhibiting smoother, differentiable curvature, whereas sparsemax is piecewise linear.

### 3.3 Computing the entmax mapping

Whether we want to use  $\alpha$ -entmax as an attention mapping, or  $L_\alpha$  as a loss function, we must be able to efficiently compute  $\mathbf{p}^* = \alpha\text{-entmax}(\mathbf{z})$ , *i.e.*, to solve the maximization in Eq. 10. For  $\alpha = 1$ , the closed-form solution is given by Eq. 4. For  $\alpha > 1$ , given  $\mathbf{z}$ , we show that there is a unique threshold  $\tau$  such that (Appendix C.1, Lemma 2):

$$\alpha\text{-entmax}(\mathbf{z}) = [(\alpha - 1)\mathbf{z} - \tau \mathbf{1}]_+^{1/\alpha - 1}, \quad (13)$$

*i.e.*, entries with score  $z_j \leq \tau/\alpha - 1$  get zero probability. For sparsemax ( $\alpha = 2$ ), the problem amounts to Euclidean projection onto  $\Delta^d$ , for which two types of algorithms are well studied:

- i. exact, based on sorting (Held et al., 1974; Michelot, 1986),
- ii. iterative, bisection-based (Liu and Ye, 2009).

The bisection approach searches for the optimal threshold  $\tau$  numerically. Blondel et al. (2019) generalize this approach in a way applicable to  $\alpha$ -entmax. The resulting algorithm is (*cf.* Appendix C.1 for details):

---

#### Algorithm 1 Compute $\alpha$ -entmax by bisection.

---

```

1 Define  $\mathbf{p}(\tau) := [\mathbf{z} - \tau \mathbf{1}]_+^{1/\alpha - 1}$ , set  $\mathbf{z} \leftarrow (\alpha - 1)\mathbf{z}$ 
2  $\tau_{\min} \leftarrow \max(\mathbf{z}) - 1$ ;  $\tau_{\max} \leftarrow \max(\mathbf{z}) - d^{1-\alpha}$ 
3 for  $t \in 1, \dots, T$  do
4    $\tau \leftarrow (\tau_{\min} + \tau_{\max})/2$ 
5    $Z \leftarrow \sum_j p_j(\tau)$ 
6   if  $Z < 1$  then  $\tau_{\max} \leftarrow \tau$  else  $\tau_{\min} \leftarrow \tau$ 
7 return  $1/Z \mathbf{p}(\tau)$ 

```

---

Algorithm 1 works by iteratively narrowing the interval containing the exact solution by exactly half. Line 7 ensures that approximate solutions are valid probability distributions, *i.e.*, that  $\mathbf{p}^* \in \Delta^d$ .

Although bisection is simple and effective, an exact sorting-based algorithm, like for sparsemax, has the potential to be faster and more accurate. Moreover, as pointed out by Condat (2016), when exact solutions are required, it is possible to construct inputs  $\mathbf{z}$  for which bisection requires arbitrarily many iterations. To address these issues, we propose a

**novel, exact algorithm for 1.5-entmax**, halfway between softmax and sparsemax.

---

**Algorithm 2** Compute  $1.5\text{-entmax}(\mathbf{z})$  exactly.

---

```

1 Sort  $\mathbf{z}$ , yielding  $z_{[d]} \leq \dots \leq z_{[1]}$ ; set  $\mathbf{z} \leftarrow \mathbf{z}/2$ 
2 for  $\rho \in 1, \dots, d$  do
3    $M(\rho) \leftarrow 1/\rho \sum_{j=1}^{\rho} z_{[j]}$ 
4    $S(\rho) \leftarrow \sum_{j=1}^{\rho} (z_{[j]} - M(\rho))^2$ 
5    $\tau(\rho) \leftarrow M(\rho) - \sqrt{1/\rho(1 - S(\rho))}$ 
6   if  $z_{[\rho+1]} \leq \tau(\rho) \leq z_{[\rho]}$  then
7     return  $\mathbf{p}^* = \lfloor \mathbf{z} - \tau \mathbf{1} \rfloor_+^2$ 
```

---

We give a full derivation in Appendix C.2. As written, Algorithm 2 is  $O(d \log d)$  because of the sort; however, in practice, when the solution  $\mathbf{p}^*$  has no more than  $k$  nonzeros, we do not need to fully sort  $\mathbf{z}$ , just to find the  $k$  largest values. Our experiments in §4.2 reveal that a partial sorting approach can be very efficient and competitive with softmax on the GPU, even for large  $d$ . Further speed-ups might be available following the strategy of Condat (2016), but our simple incremental method is very easy to implement on the GPU using primitives available in popular libraries (Paszke et al., 2017).

Our algorithm resembles the aforementioned sorting-based algorithm for projecting onto the simplex (Michelot, 1986). Both algorithms rely on the optimality conditions implying an analytically-solvable equation in  $\tau$ : for sparsemax ( $\alpha = 2$ ), this equation is linear, for  $\alpha = 1.5$  it is quadratic (Eq. 48 in Appendix C.2). Thus, exact algorithms may not be available for general values of  $\alpha$ .

### 3.4 Gradient of the entmax mapping

The following result shows how to compute the backward pass through  $\alpha$ -entmax, a requirement when using  $\alpha$ -entmax as an attention mechanism.

**Proposition 1.** *Let  $\alpha \geq 1$ . Assume we have computed  $\mathbf{p}^* = \alpha\text{-entmax}(\mathbf{z})$ , and define the vector*

$$s_i = \begin{cases} (p_i^*)^{2-\alpha}, & p_i^* > 0, \\ 0, & \text{otherwise.} \end{cases}$$

$$\text{Then, } \frac{\partial \alpha\text{-entmax}(\mathbf{z})}{\partial \mathbf{z}} = \text{diag}(\mathbf{s}) - \frac{1}{\|\mathbf{s}\|_1} \mathbf{s} \mathbf{s}^\top.$$

*Proof:* The result follows directly from the more general Proposition 2, which we state and prove in Appendix B, noting that  $\left(\frac{t^\alpha - t}{\alpha(\alpha-1)}\right)'' = t^{\alpha-2}$ . ■

The gradient expression recovers the softmax and sparsemax Jacobians with  $\alpha = 1$  and  $\alpha = 2$ ,

respectively (Martins and Astudillo, 2016, Eqs. 8 and 12), thereby providing another relationship between the two mappings. Perhaps more interestingly, Proposition 1 shows why **the sparsemax Jacobian depends only on the support** and not on the actual values of  $\mathbf{p}^*$ : the sparsemax Jacobian is equal for  $\mathbf{p}^* = [.99, .01, 0]$  and  $\mathbf{p}^* = [.5, .5, 0]$ . This is not the case for  $\alpha$ -entmax with  $\alpha \neq 2$ , suggesting that the gradients obtained with other values of  $\alpha$  may be more informative. Finally, we point out that the gradient of entmax losses involves the entmax mapping (Eq. 12), and therefore Proposition 1 also gives the Hessian of the entmax loss.

## 4 Experiments

The previous section establishes the computational building blocks required to train models with entmax sparse attention and loss functions. We now put them to use for two important NLP tasks, morphological inflection and machine translation. These two tasks highlight the characteristics of our innovations in different ways. Morphological inflection is a **character-level** task with mostly monotonic alignments, but the evaluation demands exactness: the predicted sequence must match the gold standard. On the other hand, machine translation uses a **word-level** vocabulary orders of magnitude larger and forces a sparse output layer to confront more ambiguity: any sentence has several valid translations and it is not clear beforehand that entmax will manage this well.

Despite the differences between the tasks, we keep the architecture and training procedure as similar as possible. We use two layers for encoder and decoder LSTMs and apply dropout with probability 0.3. We train with Adam (Kingma and Ba, 2015), with a base learning rate of 0.001, halved whenever the loss increases on the validation set. We use a batch size of 64. At test time, we select the model with the best validation accuracy and decode with a beam size of 5. We implemented all models with OpenNMT-py (Klein et al., 2017).<sup>2</sup>

In our primary experiments, we use three  $\alpha$  values for the attention and loss functions:  $\alpha = 1$  (softmax),  $\alpha = 1.5$  (to which our novel Algorithm 2 applies), and  $\alpha = 2$  (sparsemax). We also investigate the effect of tuning  $\alpha$  with increased granularity.

---

<sup>2</sup>Our experiment code is at <https://github.com/deep-spin/OpenNMT-entmax>.

output	$\alpha$ attention	high		medium	
		(avg.)	(ens.)	(avg.)	(ens.)
1	1	93.15	94.20	82.55	85.68
	1.5	92.32	93.50	83.20	85.63
	2	90.98	92.60	83.13	85.65
1.5	1	94.36	94.96	84.88	86.38
	1.5	94.44	95.00	84.93	86.55
	2	94.05	94.74	84.93	86.59
2	1	<b>94.59</b>	<b>95.10</b>	84.95	86.41
	1.5	94.47	95.01	<b>85.03</b>	<b>86.61</b>
	2	94.32	94.89	84.96	86.47
UZH (2018)		96.00		86.64	

Table 1: Average per-language accuracy on the test set (CoNLL–SIGMORPHON 2018 task 1) averaged or ensemble over three runs.

#### 4.1 Morphological Inflection

The goal of morphological inflection is to produce an inflected word form (such as “drawn”) given a lemma (“draw”) and a set of morphological tags ({verb, past, participle}). We use the data from task 1 of the CoNLL–SIGMORPHON 2018 shared task (Cotterell et al., 2018). shared task

**Training.** We train models under two data settings: *high* (approximately 10,000 samples per language in 86 languages) and *medium* (approximately 1,000 training samples per language in 102 languages). We depart from previous work by using **multilingual training**: each model is trained on the data from all languages in its data setting. This allows parameters to be shared between languages, eliminates the need to train language-specific models, and may provide benefits similar to other forms of data augmentation (Bergmanis et al., 2017). Each sample is presented as a pair: the source contains the lemma concatenated to the morphological tags and a special language identification token (Johnson et al., 2017; Peters et al., 2017), and the target contains the inflected form. As an example, the source sequence for Figure 1 is english\_verb\_participle\_past\_draw. Although the set of inflectional tags is not sequential, treating it as such is simple to implement and works well in practice (Kann and Schütze, 2016). All models use embedding and hidden state sizes of 300. We validate at the end of every epoch in the *high* setting and only once every ten epochs in *medium* because of its smaller size.

**Accuracy.** Results are shown in Table 1. We report the official metric of the shared task, word

accuracy averaged across languages. In addition to the average results of three individual model runs, we use an ensemble of those models, where we decode by averaging the raw probabilities at each time step. Our best sparse loss models beat the softmax baseline by nearly a full percentage point with ensembling, and up to two and a half points in the medium setting without ensembling. The choice of attention has a smaller impact. In both data settings, our best model on the validation set outperforms all submissions from the 2018 shared task except for UZH (Makarov and Clematide, 2018), which uses a more involved imitation learning approach and larger ensembles. In contrast, our only departure from standard *seq2seq* training is the drop-in replacement of softmax by entmax.

**Sparsity.** Besides their accuracy, we observed that entmax models made very sparse predictions: the best configuration in Table 1 concentrates all probability mass into a single predicted sequence in 81% validation samples in the *high* data setting, and 66% in the more difficult *medium* setting. When the model *does* give probability mass to more than one sequence, the predictions reflect reasonable ambiguity, as shown in Figure 1. Besides enhancing interpretability, sparsity in the output also has attractive properties for beam search decoding: when the beam covers all nonzero-probability hypotheses, we have a **certificate** of globally optimal decoding, **rendering beam search exact**. This is the case on 87% of validation set sequences in the *high* setting, and 79% in *medium*. To our knowledge, this is the first instance of a neural *seq2seq* model that can offer optimality guarantees.

#### 4.2 Machine Translation

We now turn to a highly different *seq2seq* regime in which the vocabulary size is much larger, there is a great deal of ambiguity, and sequences can generally be translated in several ways. We train models for three language pairs in both directions:

- IWSLT 2017 German  $\leftrightarrow$  English (DE $\leftrightarrow$ EN, Cettolo et al., 2017): training size 206,112.
- KFTT Japanese  $\leftrightarrow$  English (JA $\leftrightarrow$ EN, Neubig, 2011): training size of 329,882.
- WMT 2016 Romanian  $\leftrightarrow$  English (RO $\leftrightarrow$ EN, Bojar et al., 2016): training size 612,422, diacritics removed (following Sennrich et al., 2016b).

method	DE→EN	EN→DE	JA→EN	EN→JA	RO→EN	EN→RO
softmax	25.70 ± 0.15	21.86 ± 0.09	20.22 ± 0.08	25.21 ± 0.29	29.12 ± 0.18	28.12 ± 0.18
1.5-entmax	<b>26.17</b> ± 0.13	<b>22.42</b> ± 0.08	<b>20.55</b> ± 0.30	<b>26.00</b> ± 0.31	<b>30.15</b> ± 0.06	<b>28.84</b> ± 0.10
sparsemax	24.69 ± 0.22	20.82 ± 0.19	18.54 ± 0.11	23.84 ± 0.37	29.20 ± 0.16	28.03 ± 0.16

Table 2: Machine translation comparison of softmax, sparsemax, and the proposed 1.5-entmax as both attention mapping and loss function. Reported is tokenized test BLEU averaged across three runs (higher is better).

**Training.** We use byte pair encoding (BPE; [Sennrich et al., 2016a](#)) to ensure an open vocabulary. We use separate segmentations with 25k merge operations per language for RO↔EN and a joint segmentation with 32k merges for the other language pairs. DE↔EN is validated once every 5k steps because of its smaller size, while the other sets are validated once every 10k steps. We set the maximum number of training steps at 120k for RO↔EN and 100k for other language pairs. We use 500 dimensions for word vectors and hidden states.

**Evaluation.** Table 2 shows BLEU scores ([Papineni et al., 2002](#)) for the three models with  $\alpha \in \{1, 1.5, 2\}$ , using the same value of  $\alpha$  for the attention mechanism and loss function. We observe that the 1.5-entmax configuration consistently performs best across all six choices of language pair and direction. These results support the notion that the optimal function is somewhere between softmax and sparsemax, which motivates a more fine-grained search for  $\alpha$ ; we explore this next.

**Fine-grained impact of  $\alpha$ .** Algorithm 1 allows us to further investigate the marginal effect of varying the attention  $\alpha$  and the loss  $\alpha$ , while keeping the other fixed. We report DE→EN validation accuracy on a fine-grained  $\alpha$  grid in Figure 4. On this dataset, moving from softmax toward sparser **attention** (left) has a very small positive effect on accuracy, suggesting that the benefit in interpretability does not hurt accuracy. The impact of the **loss function**  $\alpha$  (right) is much more visible: there is a distinct optimal value around  $\alpha = 1.33$ , with performance decreasing for too large values. Interpolating between softmax and sparsemax thus inherits the benefits of both, and our novel Algorithm 2 for  $\alpha = 1.5$  is confirmed to strike a good middle ground. This experiment also confirms that bisection is effective in practice, despite being inexact. Extrapolating beyond the sparsemax loss ( $\alpha > 2$ ) does not seem to perform well.

**Sparsity.** In order to form a clearer idea of how sparse entmax becomes, we measure the average

method	# attended	# target words
softmax	24.25	17993
1.5-entmax	5.55	16.13
sparsemax	3.75	7.55

Table 3: Average number of nonzeros in the attention and output distributions for the DE→EN validation set.

number of nonzero indices on the DE→EN validation set and show it in Table 3. As expected, 1.5-entmax is less sparse than sparsemax as both an attention mechanism and output layer. In the attention mechanism, 1.5-entmax’s increased support size does not come at the cost of much interpretability, as Figure 5 demonstrates. In the output layer, 1.5-entmax assigns positive probability to only 16.13 target types out of a vocabulary of 17,993, meaning that the supported set of words often has an intuitive interpretation. Figure 2 shows the sparsity of the 1.5-entmax output layer in practice: the support becomes completely concentrated when generating a phrase like “the tree of life”, but grows when presenting a list of synonyms (“view”, “look”, “glimpse”, and so on). This has potential practical applications as a **predictive translation** system ([Green et al., 2014](#)), where the model’s support set serves as a list of candidate auto-completions at each time step.

**Training time.** Importantly, the benefits of sparsity do **not** come at a high computational cost. Our proposed Algorithm 2 for 1.5-entmax runs on the GPU at near-softmax speeds (Figure 6). For other  $\alpha$  values, bisection (Algorithm 1) is slightly more costly, but practical even for large vocabulary sizes. On DE→EN, bisection is capable of processing about 10,500 target words per second on a single Nvidia GeForce GTX 1080 GPU, compared to 13,000 words per second for 1.5-entmax with Algorithm 2 and 14,500 words per second with softmax. On the smaller-vocabulary morphology datasets, Algorithm 2 is nearly as fast as softmax.




Figure 4: Effect of tuning  $\alpha$  on DE $\rightarrow$ EN, for attention (left) and for output (right), while keeping the other  $\alpha = 1.5$ .




Figure 5: Attention weights produced by the DE $\rightarrow$ EN 1.5-entmax model. Nonzero weights are outlined.

## 5 Related Work

**Sparse attention.** Sparsity in the attention and in the output have different, but related, motivations. Sparse attention can be justified as a form of inductive bias, since for tasks such as machine translation one expects only a few source words to be relevant for each translated word. Dense attention probabilities are particularly harmful for long sequences, as shown by Luong et al. (2015), who propose “local attention” to mitigate this problem. Combining sparse attention with fertility constraints has been recently proposed by Malaviya et al. (2018). Hard attention (Xu et al., 2015; Aharoni and Goldberg, 2017; Wu et al., 2018) selects exactly one source token. Its discrete, non-differentiable nature requires imitation learning or Monte Carlo policy gradient approximations, which drastically complicate training. In contrast, entmax is a differentiable, easy to use, drop-in softmax replacement. A recent study by Jain and Wallace (2019) tackles the limitations of attention probabilities to provide interpretability. They only study dense attention in classification tasks, where attention is less crucial for the final predictions. In their conclusions, the authors defer




Figure 6: Training timing on three DE $\rightarrow$ EN runs. Markers show validation checkpoints for one of the runs.

to future work exploring sparse attention mechanisms and *seq2seq* models. We believe our paper can foster interesting investigation in this area.

**Losses for *seq2seq* models.** Mostly motivated by the challenges of large vocabulary sizes in *seq2seq*, an important research direction tackles replacing the cross-entropy loss with other losses or approximations (Bengio and Senécal, 2008; Morin and Bengio, 2005; Kumar and Tsvetkov, 2019). While differently motivated, some of the above strategies (e.g., hierarchical prediction) could be combined with our proposed sparse losses. Niculae et al. (2018) use sparsity to predict interpretable sets of structures. Since auto-regressive *seq2seq* makes no factorization assumptions, their strategy cannot be applied without approximations, such as in Edunov et al. (2018).

## 6 Conclusion and Future Work

We proposed sparse sequence-to-sequence models and provided fast algorithms to compute their attention and output transformations. Our approach yielded consistent improvements over dense models on morphological inflection and machine translation, while inducing interpretability in both attention and output distributions. Sparse output layers also provide exactness when the number of possible hypotheses does not exhaust beam search.

Given the ubiquity of softmax in NLP, entmax has many potential applications. A natural next step is to apply entmax to self-attention (Vaswani et al.,



2017). In a different vein, the strong morphological inflection results point to usefulness in other tasks where probability is concentrated in a small number of hypotheses, such as speech recognition.

## Acknowledgments

This work was supported by the European Research Council (ERC StG DeepSPIN 758969), and by the Fundação para a Ciência e Tecnologia through contracts UID/EEA/50008/2019 and CMUPERI/TIC/0046/2014 (GoLocal). We thank Mathieu Blondel, Nikolay Bogoychev, Gonçalo Correia, Erick Fonseca, Pedro Martins, Tsvetomila Mihaylova, Miguel Rios, Marcos Treviso, and the anonymous reviewers, for helpful discussion and feedback.

## References

- Roei Aharoni and Yoav Goldberg. 2017. [Morphological inflection generation with hard monotonic attention](#). In *Proc. ACL*.
- Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. 2015. [Neural machine translation by jointly learning to align and translate](#). In *Proc. ICLR*.
- Yoshua Bengio and Jean-Sébastien Senécal. 2008. [Adaptive importance sampling to accelerate training of a neural probabilistic language model](#). *IEEE Transactions on Neural Networks*, 19(4):713–722.
- Toms Bergmanis, Katharina Kann, Hinrich Schütze, and Sharon Goldwater. 2017. [Training data augmentation for low-resource morphological inflection](#). In *Proc. CoNLL–SIGMORPHON*.
- Mathieu Blondel, André FT Martins, and Vlad Niculae. 2019. [Learning classifiers with Fenchel-Young losses: Generalized entropies, margins, and algorithms](#). In *Proc. AISTATS*.
- Ondrej Bojar, Rajen Chatterjee, Christian Federmann, Yvette Graham, Barry Haddow, Matthias Huck, Antonio Jimeno Yepes, Philipp Koehn, Varvara Logacheva, Christof Monz, et al. 2016. [Findings of the 2016 conference on machine translation](#). In *ACL WMT*.
- John S Bridle. 1990. [Probabilistic interpretation of feedforward classification network outputs, with relationships to statistical pattern recognition](#). In *Neurocomputing*, pages 227–236. Springer.
- M Cettolo, M Federico, L Bentivogli, J Niehues, S Stüker, K Sudoh, K Yoshino, and C Federmann. 2017. [Overview of the IWSLT 2017 evaluation campaign](#). In *Proc. IWSLT*.
- Sumit Chopra, Michael Auli, and Alexander M Rush. 2016. [Abstractive sentence summarization with attentive recurrent neural networks](#). In *Proc. NAACL–HLT*.
- Jan K Chorowski, Dzmitry Bahdanau, Dzmitry Serdyuk, Kyunghyun Cho, and Yoshua Bengio. 2015. [Attention-based models for speech recognition](#). In *Proc. NeurIPS*.
- Laurent Condat. 2016. [Fast projection onto the simplex and the  \$\ell\_1\$  ball](#). *Mathematical Programming*, 158(1–2):575–585.
- Ryan Cotterell, Christo Kirov, John Sylak-Glassman, Géraldine Walther, Ekaterina Vylomova, Arya D McCarthy, Katharina Kann, Sebastian Mielke, Garrett Nicolai, Miikka Silfverberg, et al. 2018. [The CoNLL–SIGMORPHON 2018 shared task: Universal morphological reinflection](#). *Proc. CoNLL–SIGMORPHON*.
- John M Danskin. 1966. [The theory of max-min, with applications](#). *SIAM Journal on Applied Mathematics*, 14(4):641–664.
- Sergey Edunov, Myle Ott, Michael Auli, David Grangier, et al. 2018. [Classical structured prediction losses for sequence to sequence learning](#). In *Proc. NAACL–HLT*.
- Spence Green, Sida I Wang, Jason Chuang, Jeffrey Heer, Sebastian Schuster, and Christopher D Manning. 2014. [Human effort and machine learnability in computer aided translation](#). In *Proc. EMNLP*.
- Michael Held, Philip Wolfe, and Harlan P Crowder. 1974. [Validation of subgradient optimization](#). *Mathematical Programming*, 6(1):62–88.
- Sepp Hochreiter and Jürgen Schmidhuber. 1997. [Long short-term memory](#). *Neural Computation*, 9(8):1735–1780.
- Sarthak Jain and Byron C. Wallace. 2019. [Attention is not explanation](#). In *Proc. NAACL–HLT*.
- Melvin Johnson, Mike Schuster, Quoc V Le, Maxim Krikun, Yonghui Wu, Zhifeng Chen, Nikhil Thorat, Fernanda Viégas, Martin Wattenberg, Greg Corrado, et al. 2017. [Google’s multilingual neural machine translation system: Enabling zero-shot translation](#). *Transactions of the Association for Computational Linguistics*, 5:339–351.
- Katharina Kann and Hinrich Schütze. 2016. [MED: The LMU system for the SIGMORPHON 2016 shared task on morphological reinflection](#). In *Proc. SIGMORPHON*.
- Diederik Kingma and Jimmy Ba. 2015. [Adam: A method for stochastic optimization](#). In *Proc. ICLR*.
- Guillaume Klein, Yoon Kim, Yuntian Deng, Jean Senellart, and Alexander M Rush. 2017. [OpenNMT: Open-source toolkit for neural machine translation](#). *arXiv e-prints*.

- Sachin Kumar and Yulia Tsvetkov. 2019. Von Mises-Fisher loss for training sequence to sequence models with continuous outputs. In *Proc. ICLR*.
- Jun Liu and Jieping Ye. 2009. Efficient Euclidean projections in linear time. In *Proc. ICML*.
- Minh-Thang Luong, Hieu Pham, and Christopher D Manning. 2015. Effective approaches to attention-based neural machine translation. In *Proc. EMNLP*.
- Peter Makarov and Simon Clematide. 2018. UZH at CoNLL–SIGMORPHON 2018 shared task on universal morphological reinflection. *Proc. CoNLL–SIGMORPHON*.
- Chaitanya Malaviya, Pedro Ferreira, and André FT Martins. 2018. Sparse and constrained attention for neural machine translation. In *Proc. ACL*.
- André FT Martins and Ramón Fernandez Astudillo. 2016. From softmax to sparsemax: A sparse model of attention and multi-label classification. In *Proc. of ICML*.
- Christian Michelot. 1986. A finite algorithm for finding the projection of a point onto the canonical simplex of  $\mathbb{R}^n$ . *Journal of Optimization Theory and Applications*, 50(1):195–200.
- Frederic Morin and Yoshua Bengio. 2005. Hierarchical probabilistic neural network language model. In *Proc. AISTATS*.
- Graham Neubig. 2011. The Kyoto free translation task. <http://www.phontron.com/kftt>.
- Vlad Niculae and Mathieu Blondel. 2017. A regularized framework for sparse and structured neural attention. In *Proc. NeurIPS*.
- Vlad Niculae, André FT Martins, Mathieu Blondel, and Claire Cardie. 2018. SparseMAP: Differentiable sparse structured inference. In *Proc. ICML*.
- Kishore Papineni, Salim Roukos, Todd Ward, and Wei-Jing Zhu. 2002. BLEU: a method for automatic evaluation of machine translation. In *Proc. ACL*.
- Adam Paszke, Sam Gross, Soumith Chintala, Gregory Chanan, Edward Yang, Zachary DeVito, Zeming Lin, Alban Desmaison, Luca Antiga, and Adam Lerer. 2017. Automatic differentiation in PyTorch. In *Proc. NeurIPS Autodiff Workshop*.
- Ben Peters, Jon Dehdari, and Josef van Genabith. 2017. Massively multilingual neural grapheme-to-phoneme conversion. In *Proc. Workshop on Building Linguistically Generalizable NLP Systems*.
- R Tyrrell Rockafellar. 1970. *Convex Analysis*. Princeton University Press.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016a. Neural machine translation of rare words with subword units. In *Proc. ACL*.
- Rico Sennrich, Barry Haddow, and Alexandra Birch. 2016b. Edinburgh neural machine translation systems for WMT 16. In *Proc. WMT*.
- Constantino Tsallis. 1988. Possible generalization of Boltzmann-Gibbs statistics. *Journal of Statistical Physics*, 52:479–487.
- Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is all you need. In *Proc. NeurIPS*.
- Martin J Wainwright and Michael I Jordan. 2008. Graphical models, exponential families, and variational inference. *Foundations and Trends® in Machine Learning*, 1(1–2):1–305.
- Shijie Wu, Pamela Shapiro, and Ryan Cotterell. 2018. Hard non-monotonic attention for character-level transduction. In *Proc. EMNLP*.
- Kelvin Xu, Jimmy Ba, Ryan Kiros, Kyunghyun Cho, Aaron Courville, Ruslan Salakhudinov, Rich Zemel, and Yoshua Bengio. 2015. Show, attend and tell: Neural image caption generation with visual attention. In *Proc. ICML*.

## A Background

### A.1 Tsallis entropies

Recall the definition of the Tsallis family of entropies in Eq. 9 for  $\alpha \geq 1$ ,

$$H_\alpha^T(\mathbf{p}) := \begin{cases} \frac{1}{\alpha(\alpha-1)} \sum_j (p_j - p_j^\alpha), & \text{if } \alpha \neq 1, \\ H^S(\mathbf{p}), & \text{if } \alpha = 1. \end{cases} \quad (14)$$

This family is **continuous** in  $\alpha$ , *i.e.*,

$$\lim_{\alpha \rightarrow 1} H_\alpha^T(\mathbf{p}) = H_1^T(\mathbf{p})$$

for any  $\mathbf{p} \in \Delta^d$ .

*Proof:* We rewrite  $H_\alpha^T$  in separable form:

$$H_\alpha^T(\mathbf{p}) = \sum_j h_\alpha(p_j), \quad \text{with} \quad (15)$$

$$h_\alpha(t) := \begin{cases} \frac{t-t^\alpha}{\alpha(\alpha-1)}, & \alpha > 1, \\ -t \log t, & \alpha = 1. \end{cases}$$

It suffices to show that  $\lim_{\alpha \rightarrow 1} h_\alpha(t) = h_1(t)$  for  $t \in [0, 1]$ . Let  $f(\alpha) := t - t^\alpha$ , and  $g(\alpha) := \alpha(\alpha - 1)$ . Note that  $\frac{f(1)}{g(1)} = 0/0$ , leading to an indetermination. Take the derivatives of  $f$  and  $g$ :

$$\begin{aligned} f'(\alpha) &= 0 - (\exp(\log t^\alpha))' \\ &= -\exp(\log t^\alpha) \cdot (\alpha \log t)' \\ &= -t^\alpha \log t, \end{aligned} \quad (16)$$

and  $g'(\alpha) = 2\alpha - 1$ . From l'Hôpital's rule,

$$\lim_{\alpha \rightarrow 1} \frac{f(\alpha)}{g(\alpha)} = \lim_{\alpha \rightarrow 1} \frac{f'(\alpha)}{g'(\alpha)} = -t \log t = h_1(t).$$

Note also that, as  $\alpha \rightarrow \infty$ , the denominator grows unbounded, so  $H_\infty^T \equiv 0$ . ■

### A.2 Fenchel-Young losses

In this section, we recall the definitions and properties essential for our construction of  $\alpha$ -entmax. The concepts below were formalized by Blondel et al. (2019) in more generality; we present below a less general version, sufficient for our needs.

**Definition 1** (Probabilistic prediction function regularized by  $\Omega$ ). *Let  $\Omega: \Delta^d \rightarrow \mathbb{R} \cup \{\infty\}$  be a strictly convex regularization function. We define the prediction function  $\pi_\Omega$  as*

$$\pi_\Omega(\mathbf{z}) \in \operatorname{argmax}_{\mathbf{p} \in \Delta^d} (\mathbf{p}^\top \mathbf{z} - \Omega(\mathbf{p})) \quad (17)$$

**Definition 2** (Fenchel-Young loss generated by  $\Omega$ ). *Let  $\Omega: \Delta^d \rightarrow \mathbb{R} \cup \{\infty\}$  be a strictly convex regularization function. Let  $\mathbf{y} \in \Delta$  denote a ground-truth label (for example,  $\mathbf{y} = \mathbf{e}_y$  if there is a unique correct class  $y$ ). Denote by  $\mathbf{z} \in \mathbb{R}^d$  the prediction scores produced by some model, and by  $\mathbf{p}^* := \pi_\Omega(\mathbf{z})$  the probabilistic predictions. The **Fenchel-Young loss**  $L_\Omega: \mathbb{R}^d \times \Delta \rightarrow \mathbb{R}_+$  generated by  $\Omega$  is*

$$L_\Omega(\mathbf{z}; \mathbf{y}) := \Omega(\mathbf{y}) - \Omega(\mathbf{p}^*) + \mathbf{z}^\top (\mathbf{p}^* - \mathbf{y}). \quad (18)$$

This justifies our choice of entmax mapping and loss (Eqs. 10–11), as  $\pi_{-H_\alpha^T} = \alpha$ -entmax and  $L_{-H_\alpha^T} = L_\alpha$ .

**Properties of Fenchel-Young losses.**

1. **Non-negativity.**  $L_\Omega(\mathbf{z}; \mathbf{y}) \geq 0$  for any  $\mathbf{z} \in \mathbb{R}^d$  and  $\mathbf{y} \in \Delta^d$ .
2. **Zero loss.**  $L_\Omega(\mathbf{z}; \mathbf{y}) = 0$  if and only if  $\mathbf{y} = \pi_\Omega(\mathbf{z})$ , *i.e.*, the prediction is exactly correct.
3. **Convexity.**  $L_\Omega$  is convex in  $\mathbf{z}$ .
4. **Differentiability.**  $L_\Omega$  is differentiable with  $\nabla L_\Omega(\mathbf{z}; \mathbf{y}) = \mathbf{p}^* - \mathbf{y}$ .
5. **Smoothness.** If  $\Omega$  is strongly convex, then  $L_\Omega$  is smooth.
6. **Temperature scaling.** For any constant  $t > 0$ ,  $L_{t\Omega}(\mathbf{z}; \mathbf{y}) = tL_\Omega(\mathbf{z}/t; \mathbf{y})$ .

**Characterizing the solution  $\mathbf{p}^*$  of  $\pi_\Omega(\mathbf{z})$ .** To shed light on the generic probability mapping in Eq. 17, we derive below the optimality conditions characterizing its solution. The optimality conditions are essential not only for constructing algorithms for computing  $\mathbf{p}^*$  (Appendix C), but also for deriving the Jacobian of the mapping (Appendix B). The Lagrangian of the maximization in Eq. 17 is

$$\mathcal{L}(\mathbf{p}, \boldsymbol{\nu}, \tau) = \Omega(\mathbf{p}) - (\mathbf{z} + \boldsymbol{\nu})^\top \mathbf{p} + \tau(\mathbf{1}^\top \mathbf{p} - 1). \quad (19)$$

with subgradient

$$\partial_p \mathcal{L}(\mathbf{p}, \boldsymbol{\nu}, \tau) = \partial \Omega(\mathbf{p}) - \mathbf{z} - \boldsymbol{\nu} + \tau \mathbf{1}. \quad (20)$$

The subgradient KKT conditions are therefore:

$$\begin{cases} \mathbf{z} + \boldsymbol{\nu} - \tau \mathbf{1} \in \partial \Omega(\mathbf{p}) & (21) \\ \mathbf{p}^\top \boldsymbol{\nu} = 0 & (22) \\ \mathbf{p} \in \Delta^d & (23) \\ \boldsymbol{\nu} \geq 0. & (24) \end{cases}$$

**Connection to softmax and sparsemax.** We may now directly see that, when  $\Omega(\mathbf{p}) := \sum_j p_j \log p_j$ , Eq. 21 becomes  $\log p_j = z_j + \nu_j - \tau - 1$ , which can only be satisfied if  $p_j > 0$ , thus  $\boldsymbol{\nu} = \mathbf{0}$ . Then,  $p_j = \exp(z_j)/Z$ , where  $Z := \exp(\tau + 1)$ . From Eq. 23,  $Z$  must be such that  $p_j$  sums to 1, yielding the well-known softmax expression. In the case of sparsemax, note that for any  $\mathbf{p} \in \Delta^d$ , we have

$$\begin{aligned} \Omega(\mathbf{p}) &= -\text{H}^G(\mathbf{p}) \\ &= 1/2 \sum_j p_j (p_j - 1) \\ &= 1/2 \|\mathbf{p}\|^2 - \frac{1}{2} \underbrace{\sum_j p_j}_{=1} \quad (25) \\ &= 1/2 \|\mathbf{p}\|^2 + \text{const.} \end{aligned}$$

Thus,

$$\begin{aligned} &\operatorname{argmax}_{\mathbf{p} \in \Delta^d} \mathbf{p}^\top \mathbf{z} + \text{H}^G(\mathbf{p}) \\ &= \operatorname{argmin}_{\mathbf{p} \in \Delta^d} 0.5 \left( \|\mathbf{p}\|^2 - 2\mathbf{p}^\top \mathbf{z} + \|\mathbf{z}\|^2 \right) \quad (26) \\ &= \operatorname{argmin}_{\mathbf{p} \in \Delta^d} \|\mathbf{p} - \mathbf{z}\|^2. \end{aligned}$$

## B Backward pass for generalized sparse attention mappings

When a mapping  $\pi_\Omega$  is used inside the computation graph of a neural network, the Jacobian of the mapping has the important role of showing how to propagate error information, necessary when training with gradient methods. In this section, we derive a new, simple expression for the Jacobian of generalized sparse mappings  $\pi_\Omega$ . We apply this result to obtain a simple form for the Jacobian of  $\alpha$ -entmax mappings.

The proof is in two steps. First, we prove a lemma that shows that Jacobians are zero outside of the support of the solution. Then, to complete the result, we characterize the Jacobian on the support.

**Lemma 1** (Sparse attention mechanisms have sparse Jacobians). *Let  $\Omega : \mathbb{R}^d \rightarrow \mathbb{R}$  be strongly convex. The attention mapping  $\pi_\Omega$  is differentiable almost everywhere, with Jacobian  $\frac{\partial \pi_\Omega}{\partial \mathbf{z}}$  symmetric and satisfying*

$$\frac{\partial(\pi_\Omega(\mathbf{z}))_i}{\partial z_j} = 0$$

if  $(\pi_\Omega(\mathbf{z}))_i = 0$  or  $(\pi_\Omega(\mathbf{z}))_j = 0$ .

*Proof:* Since  $\Omega$  is strictly convex, the argmax in Eq. 17 is unique. Using Danskin's theorem (Danskin, 1966), we may write

$$\pi_\Omega(\mathbf{z}) = \nabla \max_{\mathbf{p} \in \Delta} (\mathbf{p}^\top \mathbf{z} - \Omega(\mathbf{p})) = \nabla \Omega^*(\mathbf{z}).$$

Since  $\Omega$  is strongly convex, the gradient of its conjugate  $\Omega^*$  is differentiable almost everywhere (Rockafellar, 1970). Moreover,  $\frac{\partial \pi_\Omega}{\partial \mathbf{z}}$  is the Hessian of  $\Omega^*$ , therefore symmetric, proving the first two claims.

Recall the definition of a partial derivative,

$$\frac{\partial(\pi_\Omega(\mathbf{z}))_i}{\partial z_j} = \lim_{\varepsilon \rightarrow 0} \frac{1}{\varepsilon} (\pi_\Omega(\mathbf{z} + \varepsilon \mathbf{e}_j)_i - \pi_\Omega(\mathbf{z})_i).$$

Denote by  $\mathbf{p}^* := \pi_\Omega(\mathbf{z})$ . We will show that for any  $j$  such that  $p_j^* = 0$ , and any  $\varepsilon \geq 0$ ,

$$\pi_\Omega(\mathbf{z} - \varepsilon \mathbf{e}_j) = \pi_\Omega(\mathbf{z}) = \mathbf{p}^*.$$

In other words, we consider only one side of the limit, namely **subtracting** a small non-negative  $\varepsilon$ . A vector  $\mathbf{p}^*$  solves the optimization problem in Eq. 17 if and only if there exists  $\boldsymbol{\nu}^* \in \mathbb{R}^d$  and  $\tau^* \in \mathbb{R}$  satisfying Eqs. 21–24. Let  $\boldsymbol{\nu}_\varepsilon := \boldsymbol{\nu}^* + \varepsilon \mathbf{e}_j$ . We verify that  $(\mathbf{p}^*, \boldsymbol{\nu}_\varepsilon, \tau^*)$  satisfies the optimality conditions for  $\pi_\Omega(\mathbf{z} - \varepsilon \mathbf{e}_j)$ , which implies that  $\pi_\Omega(\mathbf{z} - \varepsilon \mathbf{e}_j) = \pi_\Omega(\mathbf{z})$ . Since we add a non-negative quantity to  $\boldsymbol{\nu}^*$ , which is non-negative to begin with,  $(\boldsymbol{\nu}_\varepsilon)_j \geq 0$ , and since  $p_j^* = 0$ , we also satisfy  $p_j^*(\boldsymbol{\nu}_\varepsilon)_j = 0$ . Finally,

$$\begin{aligned} &\mathbf{z} - \varepsilon \mathbf{e}_j + \boldsymbol{\nu}_\varepsilon - \tau^* \mathbf{1} \\ &= \mathbf{z} - \varepsilon \mathbf{e}_j + \boldsymbol{\nu}^* + \varepsilon \mathbf{e}_j - \tau^* \mathbf{1} \quad (27) \\ &\in \partial \Omega(\mathbf{p}^*). \end{aligned}$$

It follows that

$$\lim_{\varepsilon \rightarrow 0^-} \frac{1}{\varepsilon} (\pi_\Omega(\mathbf{z} + \varepsilon \mathbf{e}_j)_i - \pi_\Omega(\mathbf{z})_i) = 0. \quad (28)$$

If  $\pi_\Omega$  is differentiable at  $\mathbf{z}$ , this one-sided limit must agree with the derivative. Otherwise, the sparse one-sided limit is a generalized Jacobian. ■

**Proposition 2.** *Let  $\mathbf{p}^* := \pi_\Omega(\mathbf{z})$ , with strongly convex and differentiable  $\Omega$ . Denote the support of  $\mathbf{p}^*$  by  $\mathcal{S} = \{j \in \{1, \dots, d\} : p_j > 0\}$ . If the second derivative  $h_{ij} = \frac{\partial^2 \Omega}{\partial p_i \partial p_j}(\mathbf{p}^*)$  exists for any  $i, j \in \mathcal{S}$ , then*

$$\frac{\partial \pi_\Omega}{\partial \mathbf{z}} = \mathbf{S} - \frac{1}{\|\mathbf{s}\|_1} \mathbf{s} \mathbf{s}^\top \quad (29)$$

where

$$\mathbf{S}_{ij} = \begin{cases} \mathbf{H}_{ij}^{-1}, & i, j \in \mathcal{S}, \\ 0, & \text{o.w.} \end{cases}, \text{ and } \mathbf{s} = \mathbf{S}\mathbf{1}. \quad (30)$$

In particular, if  $\Omega(\mathbf{p}) = \sum_j g(p_j)$  with  $g$  twice differentiable on  $(0, 1]$ , we have

$$\frac{\partial \pi_\Omega}{\partial \mathbf{z}} = \text{diag } \mathbf{s} - \frac{1}{\|\mathbf{s}\|_1} \mathbf{s} \mathbf{s}^\top \quad (31)$$

where

$$s_i = \begin{cases} (g''(p_i^*))^{-1}, & i \in \mathcal{S}, \\ 0, & \text{o.w.} \end{cases} \quad (32)$$

*Proof:* Lemma 1 verifies that  $\frac{\partial(\pi_\Omega)_i}{\partial z_j} = 0$  for  $i, j \notin \mathcal{S}$ . It remains to find the derivatives w.r.t.  $i, j \in \mathcal{S}$ . Denote by  $\bar{\mathbf{p}}^*, \bar{z}$  the restriction of the corresponding vectors to the indices in the support  $\mathcal{S}$ . The optimality conditions on the support are

$$\begin{cases} \mathbf{g}(\bar{\mathbf{p}}) + \tau \mathbf{1} &= \bar{z} \\ \mathbf{1}^\top \bar{\mathbf{p}} &= 1 \end{cases} \quad (33)$$

where  $\mathbf{g}(\bar{\mathbf{p}}) := (\nabla \Omega(\mathbf{p}))|_{\mathcal{S}}$ , so  $\frac{\partial \mathbf{g}}{\partial \bar{\mathbf{p}}}(\bar{\mathbf{p}}^*) = \mathbf{H}$ . Differentiating w.r.t.  $\bar{z}$  at  $\bar{\mathbf{p}}^*$  yields

$$\begin{cases} \mathbf{H} \frac{\partial \bar{\mathbf{p}}}{\partial \bar{z}} + \mathbf{1} \frac{\partial \tau}{\partial \bar{z}} &= \mathbf{I} \\ \mathbf{1}^\top \frac{\partial \bar{\mathbf{p}}}{\partial \bar{z}} &= 0 \end{cases} \quad (34)$$

Since  $\Omega$  is strictly convex,  $\mathbf{H}$  is invertible. From Gaussian elimination (*i.e.*, the Schur complement),

$$\frac{\partial \tau}{\partial \bar{z}} = -\frac{1}{\mathbf{1}^\top \mathbf{H}^{-1} \mathbf{1}} \mathbf{1}^\top \mathbf{H}^{-1},$$

which can then be used to solve for  $\frac{\partial \bar{\mathbf{p}}}{\partial \bar{z}}$  giving

$$\frac{\partial \bar{\mathbf{p}}}{\partial \bar{z}} = \mathbf{H}^{-1} - \frac{1}{\mathbf{1}^\top \mathbf{H}^{-1} \mathbf{1}} \mathbf{H}^{-1} \mathbf{1} \mathbf{1}^\top \mathbf{H}^{-1},$$

yielding the desired result. When  $\Omega$  is separable,  $\mathbf{H}$  is diagonal, with  $H_{ii} = g''(p_i^*)$ , yielding the simplified expression which completes the proof. ■

**Connection to other differentiable attention results.** Our result is similar, but simpler than [Niculae and Blondel \(2017, Proposition 1\)](#), especially in the case of separable  $\Omega$ . Crucially, our result does not require that the second derivative exist outside of the support. As such, unlike the cited work, our result is applicable in the case of  $\alpha$ -entmax, where either  $g''(t) = t^{\alpha-2}$  or its reciprocal may not exist at  $t = 0$ .

## C Algorithms for entmax

### C.1 General thresholded form for bisection algorithms.

The following lemma provides a simplified form for the solution of  $\alpha$ -entmax.

**Lemma 2.** *For any  $\mathbf{z} \in \mathbb{R}^d$ , there is a unique  $\tau^* \in \mathbb{R}$  such that*

$$\alpha\text{-entmax}(\mathbf{z}) = [(\alpha - 1)\mathbf{z} - \tau^* \mathbf{1}]_+^{1/\alpha-1}. \quad (35)$$

*Proof:* We use the regularized prediction functions defined in [Appendix A.2](#). From both definitions,

$$\alpha\text{-entmax}(\mathbf{z}) \equiv \pi_{-\mathbf{H}_\alpha^\top}(\mathbf{z}).$$

We first note that for all  $\mathbf{p} \in \Delta^d$ ,

$$-(\alpha - 1)\mathbf{H}_\alpha^\top(\mathbf{p}) = \frac{1}{\alpha} \sum_{i=1}^d p_i^\alpha + \text{const.} \quad (36)$$

From the constant invariance and scaling of  $\pi_\Omega$  ([Blondel et al., 2019, Proposition 1, items 4–5](#)),

$$\pi_{-\mathbf{H}_\alpha^\top}(\mathbf{z}) = \pi_\Omega((\alpha - 1)\mathbf{z}), \quad (37)$$

with

$$\Omega(\mathbf{p}) = \sum_{j=1}^d g(p_j), \quad g(t) = \frac{t^\alpha}{\alpha}. \quad (38)$$

Using ([Blondel et al., 2019, Proposition 5](#)), noting that  $g'(t) = t^{\alpha-1}$  and  $(g')^{-1}(u) = u^{1/\alpha-1}$ , yields

$$\pi_\Omega(\mathbf{z}) = [\mathbf{z} - \tau^* \mathbf{1}]_+^{1/\alpha-1}, \quad (39)$$

and therefore

$$\alpha\text{-entmax}(\mathbf{z}) = [(\alpha - 1)\mathbf{z} - \tau^* \mathbf{1}]_+^{1/\alpha-1}. \quad (40)$$

Uniqueness of  $\tau^*$  follows from the fact that  $\alpha$ -entmax has a unique solution  $\mathbf{p}^*$ , and [Eq. 40](#) implies a one-to-one mapping between  $\mathbf{p}^*$  and  $\tau^*$ , as long as  $\mathbf{p}^* \in \Delta$ . ■

**Corollary 2.1.** *For  $\alpha = 1.5$ , Lemma 2 implies existence of a unique  $\tau^*$  such that*

$$1.5\text{-entmax}(\mathbf{z}) = [\mathbf{z}/2 - \tau^* \mathbf{1}]_+^2.$$

## C.2 An exact algorithm for entmax with $\alpha = 1.5$ : Derivation of Algorithm 2.

In this section, we derive an exact, sorting-based algorithm for 1.5-entmax. The key observation is that the solution can be characterized by the size of its support,  $\rho^* = \|\mathbf{p}^*\|_0$ . Then, we can simply enumerate all possible values of  $\rho \in \{1, \dots, d\}$  until the solution verifies all optimality conditions. The challenge, however, is expressing the threshold  $\tau$  as a function of the support size  $\rho$ ; for this, we rely on  $\alpha = 1.5$ .

**Proposition 3.** *Exact computation of 1.5-entmax( $\mathbf{z}$ )*

Let  $z_{[d]} \leq \dots \leq z_{[1]}$  denote the sorted coordinates of  $\mathbf{z}$ , and, for convenience, let  $z_{[d+1]} := -\infty$ . Define the top- $\rho$  mean, unnormalized variance, and induced threshold for  $\rho \in \{1, \dots, d\}$  as

$$\begin{aligned} M_{\mathbf{z}}(\rho) &:= \frac{1}{\rho} \sum_{j=1}^{\rho} z_{[j]}, \\ S_{\mathbf{z}}(\rho) &:= \sum_{j=1}^{\rho} (z_{[j]} - M_{\mathbf{z}}(\rho))^2, \\ \tau_{\mathbf{z}}(\rho) &:= \begin{cases} M_{\mathbf{z}}(\rho) - \sqrt{\frac{1-S_{\mathbf{z}}(\rho)}{\rho}}, & S_{\mathbf{z}}(\rho) \leq 1, \\ +\infty, & S_{\mathbf{z}}(\rho) > 1. \end{cases} \end{aligned}$$

Then,

$$(1.5\text{-entmax}(\mathbf{z}))_i = \left[ \frac{z_i}{2} - \tau_{z/2}(\rho) \right]_+, \quad (41)$$

for any  $\rho$  satisfying  $\tau_{\mathbf{z}}(\rho) \in [z_{[\rho+1]}, z_{[\rho]}]$ .

Proposition 3 implies the correctness of Algorithm 2. To prove it, we first show the following.

**Lemma 3.** *Define  $\tau(\rho)$  as in Proposition 3. Then,  $\tau$  is non-decreasing, and there exists  $\rho_{\max} \in \{1, \dots, d\}$  such that  $\tau$  is finite for  $1 \leq \rho \leq \rho_{\max}$ , and infinite for  $\rho > \rho_{\max}$ .*

The proof is slightly more technical, and we defer it to *after* the proof of the proposition.

**Proof of Proposition 3.** First, using Corollary 2.1 we reduce the problem of computing 1.5-entmax to

$$\pi_{\Omega}(\mathbf{z}) := \operatorname{argmax}_{\mathbf{p} \in \Delta^d} \mathbf{p}^{\top} \mathbf{z} - \sum_j \frac{2}{3} p_j^{3/2}. \quad (42)$$

Denote by  $\tau^*$  the optimal threshold as defined in the corollary. We will show that  $\tau^* = \tau(\rho)$  for any  $\rho$  satisfying  $\tau(\rho) \in [z_{[\rho+1]}, z_{[\rho]}]$ , where we assume, for convenience,  $z_{[d+1]} = -\infty$ . The generic

stationarity condition in Eq. 21, applied to the problem in Eq. 42, takes the form

$$\sqrt{p_j} = \nu_j + z_j - \tau \quad \forall 0 < j \leq d \quad (43)$$

Since  $\Omega$  is symmetric,  $\pi_{\Omega}$  is permutation-preserving (Blondel et al., 2019, Proposition 1, item 1), so we may assume w.l.o.g. that  $\mathbf{z}$  is sorted non-increasingly, i.e.,  $z_1 \geq \dots \geq z_d$ ; in other words,  $z_j = z_{[j]}$ . Therefore, the optimal  $\mathbf{p}$  is also non-increasing. Denote by  $\rho$  an index such as  $p_j \geq 0$  for  $1 \leq j \leq \rho$ , and  $p_j = 0$  for  $j > \rho$ . From the complementary slackness condition (22),  $\nu_j = 0$  for  $1 \leq j \leq \rho$ , thus we may split the stationarity conditions (43) into

$$\begin{cases} \sqrt{p_j} = z_j - \tau, & \forall 1 \leq j \leq \rho, \\ \nu_j = \tau - z_j, & \forall \rho < j \leq d. \end{cases} \quad (44)$$

For (44) to have solutions, the r.h.s. must be non-negative, i.e.,  $\tau \leq z_j$  for  $j \leq \rho$ , so  $\tau \leq z_{\rho}$ . At the same time, from dual feasibility (24) we have  $\nu_j = \tau - z_j \geq 0$  for  $j > \rho$ , therefore

$$\tau(\rho) \in [z_{\rho+1}, z_{\rho}]. \quad (46)$$

Given  $\rho$ , we can solve for  $\tau$  using (44) and primal feasibility (23)

$$1 = \sum_{j=1}^d p_j = \sum_{j=1}^{\rho} (z_j - \tau)^2. \quad (47)$$

Expanding the squares and dividing by  $2\rho$  yields the quadratic equation

$$\frac{1}{2}\tau^2 - \frac{\sum_{j=1}^{\rho} z_j}{\rho}\tau + \frac{\sum_{j=1}^{\rho} z_j^2 - 1}{2\rho} = 0, \quad (48)$$

with discriminant

$$\Delta(\rho) = (M(\rho))^2 - \frac{\sum_{j=1}^{\rho} z_j^2}{\rho} + \frac{1}{\rho} = \frac{1 - S(\rho)}{\rho}. \quad (49)$$

where we used the variance expression  $\mathbb{E}[(X - \mathbb{E}[X])^2] = \mathbb{E}[X^2] - \mathbb{E}[X]^2$ . If  $S(\rho) > 1$ ,  $\Delta(\rho) < 0$ , so there must exist an optimal  $\rho$  satisfying  $S(\rho) \in [0, 1]$ . Therefore, (48) has the two solutions  $\tau_{\pm}(\rho) = M(\rho) \pm \sqrt{\frac{1-S(\rho)}{\rho}}$ . However,  $\tau_+$  leads to a contradiction: The mean  $M(\rho)$  is never smaller than the smallest averaged term, so  $M(\rho) \geq z_{\rho}$ , and thus  $\tau_+ \geq z_{\rho}$ . At the same time, from (46),  $\tau \leq z_{\rho}$ , so  $\tau$  must equal  $z_{\rho}$ , which can only happen if  $M(\rho) = z_{\rho}$  and

$S(\rho) = 1$ . But  $M(\rho) = z_\rho$  only if  $z_1 = \dots = z_\rho$ , in which case  $S(\rho) = 0$  (contradiction).

Therefore,  $\tau^* = \tau(\rho) = M(\rho) - \sqrt{\frac{1-S(\rho)}{\rho}}$  for some  $\rho$  verifying (46). It remains to show that any such  $\rho$  leads to the same value of  $\tau(\rho)$ . Pick any  $\rho_1 < \rho_2$ , both verifying (46). Therefore,  $\rho_1 + 1 \leq \rho_2$  and

$$\begin{aligned} z_{\rho_1+1} &\stackrel{(46) \text{ for } \rho_1}{\leq} \tau(\rho_1) \\ &\stackrel{\text{Lemma 3}}{\leq} \tau(\rho_2) \\ &\stackrel{(46) \text{ for } \rho_2}{\leq} z_{\rho_2} \\ &\stackrel{z \text{ sorted}}{\leq} z_{\rho_1+1}, \end{aligned} \quad (50)$$

thus  $\tau(\rho_1) = \tau(\rho_2)$ , and so any  $\rho$  verifying (46) satisfies  $\tau^* = \tau(\rho)$ , concluding the proof.  $\square$

**Proof of Lemma 3.** We regard  $\tau(\rho)$  as an extended-value sequence, *i.e.*, a function from  $\mathbb{N} \rightarrow \mathbb{R} \cup \infty$ . The lemma makes a claim about the *domain* of the sequence  $\tau$ , and a claim about its monotonicity. We prove the two in turn.

**Domain of  $\tau$ .** The threshold  $\tau(\rho)$  is only finite for  $\rho \in T := \{\rho \in \{1, \dots, d\} : S(\rho) \leq 1\}$ , *i.e.*, where  $(1-S(\rho))/\rho \geq 0$ . We show there exists  $\rho_{\max}$  such that  $T = \{1, \dots, \rho_{\max}\}$ . Choose  $\rho_{\max}$  as the largest index satisfying  $S(\rho_{\max}) \leq 1$ . By definition,  $\rho > \rho_{\max}$  implies  $\rho \notin T$ . Remark that  $S(1) = 0$ , and  $S(\rho+1) - S(\rho) = (\cdot)^2 \geq 0$ . Therefore,  $S$  is nondecreasing and, for any  $1 \leq \rho \leq \rho_{\max}$ ,  $0 \leq S(\rho) \leq 1$ .

**Monotonicity of  $\tau$ .** Fix  $\rho \in [\rho_{\max} - 1]$ , assume w.l.o.g. that  $M_z(\rho) = 0$ , and define  $\tilde{z}$  as

$$\tilde{z}_{[j]} = \begin{cases} x, & j = \rho + 1, \\ z_{[j]}, & \text{otherwise.} \end{cases}$$

The  $\rho$  highest entries of  $\tilde{z}$  are the same as in  $z$ , so  $M_{\tilde{z}}(\rho) = M_z(\rho) = 0$ ,  $S_{\tilde{z}}(\rho) = S_z(\rho)$ , and  $\tau_{\tilde{z}}(\rho) = \tau_z(\rho)$ . Denote  $\tilde{\tau}(x) := \tau_{\tilde{z}}(\rho + 1)$ , and analogously  $\tilde{M}(x)$  and  $\tilde{S}(x)$ . Then,

$$\tau_z(\rho+1) = \tilde{\tau}(z_{[\rho+1]}) \geq \min_{x: \tilde{S}(x) \in [0,1]} \tilde{\tau}(x) =: \tilde{\tau}(x^*) \quad (51)$$

We seek the lower bound  $\tilde{\tau}(x^*)$  and show that  $\tilde{\tau}(x^*) \geq \tau_z(\rho)$ . From (51), this implies  $\tau_z(\rho + 1) \geq \tau_z(\rho)$  and, by transitivity, the monotonicity of  $\tau_z$ .

It is easy to verify that the following incremental update expressions hold.

$$\tilde{M}(x) = \frac{x}{\rho+1}, \quad \tilde{S}(x) = S_z(\rho) + \frac{\rho}{\rho+1}x^2. \quad (52)$$

We must solve the optimization problem

$$\text{minimize}_x \tilde{\tau}(x) \quad \text{subject to} \quad \tilde{S}(x) \in [0, 1]. \quad (53)$$

The objective value is

$$\begin{aligned} \tilde{\tau}(x) &= \tilde{M}(x) - \sqrt{\frac{1-\tilde{S}(x)}{\rho+1}} \\ &= \frac{1}{\rho+1} \left( x - \sqrt{(1-S_z(\rho))(\rho+1) - \rho x^2} \right) \end{aligned} \quad (54)$$

Ignoring the constraint for a moment and setting the gradient to 0 yields the solution

$$\begin{aligned} 0 &= \tilde{\tau}'(x^*) \\ &= \frac{1}{\rho+1} \left( 1 + \frac{\rho x^*}{\sqrt{(1-S_z(\rho))(\rho+1) - \rho x^{*2}}} \right) \\ &\iff \rho x^* = -\sqrt{(1-S_z(\rho))(\rho+1) - \rho x^{*2}}, \end{aligned} \quad (55)$$

implying  $x^* < 0$ . Squaring both sides and rearranging yields the solution of the unconstrained optimization problem,

$$x^* = -\sqrt{\frac{1-S_z(\rho)}{\rho}}. \quad (56)$$

We verify that  $x^*$  readily satisfies the constraints, thus it is a solution to the minimization in Eq. 53. Since

$$\tilde{S}(x^*) = S_z(\rho) + \frac{1-S_z(\rho)}{\rho+1}, \quad (57)$$

we have

$$\tilde{S}(x^*) \geq S_z(\rho) \geq 0 \quad (58)$$

and

$$\tilde{S}(x^*) \leq S_z(\rho) + \frac{1-S_z(\rho)}{2} \leq 1. \quad (59)$$

Plugging  $x^*$  into the objective yields

$$\begin{aligned}\tilde{\tau}(x^*) &= \frac{1}{\rho+1} \left( -\sqrt{\frac{1-S_z(\rho)}{\rho}} - \sqrt{\rho(1-S_z(\rho))} \right) \\ &= -\sqrt{\frac{1-S_z(\rho)}{\rho}} \frac{1}{\rho+1} (1+\rho) \\ &= -\sqrt{\frac{1-S_z(\rho)}{\rho}} = \tau_z(\rho).\end{aligned}\tag{60}$$

Therefore,  $\tilde{\tau}(x) \geq \tau_z(\rho)$  for any valid  $x$ , proving that  $\tau_z(\rho) \leq \tau_z(\rho+1)$ .