

000
001
002
003
004
005
006
007
008
009
010
011
012
013
014
015
016
017
018
019
020
021
022
023
024
025
026
027
028
029
030
031
032
033
034
035
036
037
038
039
040
041
042
043
044
045
046
047
048
049
050
051
052
053

AVEY BIDIRECTIONAL ARCHITECTURE

Anonymous authors

Paper under double-blind review

ABSTRACT

Compact pretrained bidirectional encoders remain the backbone of industrial NLP under tight compute and memory budgets. Their effectiveness stems from self-attention’s ability to deliver bidirectional contextualization with high parallelism, as popularized by BERT-style architectures. Recently, Avey was introduced as an autoregressive, attention-free alternative that naturally admits an encoder-only adaptation. In this paper, we reformulate Avey for the encoder-only paradigm and propose several innovations to its architecture, including decoupled static and dynamic parameterizations, stability-oriented normalization, and neural compression. Results show that this reformulated architecture compares favorably to four widely used Transformer-based encoders, consistently outperforming them on standard token-classification and information-retrieval benchmarks while scaling more efficiently to long contexts.

1 INTRODUCTION

Pretrained bidirectional Transformer encoders, most notably BERT (Devlin et al., 2019), have been especially impactful in resource-constrained, application-specific settings, where compact models can be efficiently fine-tuned for downstream tasks and deployed under strict latency and memory budgets. Unlike unidirectional Transformer decoders, bidirectional encoders condition each token on both its left and right contexts, yielding fully contextualized representations that improve disambiguation and translate into stronger performance on certain discriminative tasks (e.g., classification, retrieval, and extractive question-answering) (Liu et al., 2019; Wang et al., 2019b;a; Karpukhin et al., 2020; Khattab & Zaharia, 2020; Rajpurkar et al., 2016; 2018). Since BERT’s introduction, such encoders have seen broad and sustained adoption across academia and industry (Muennighoff et al., 2023; Thakur et al., 2021; Santhanam et al., 2022; Wang et al., 2022; Su et al., 2023), particularly targeting high-throughput, high-precision, and budget-constrained applications (Lan et al., 2020; Sanh et al., 2019; Sun et al., 2020; Jiao et al., 2020).

The BERT family’s success in research and industry was enabled by the Transformer (Vaswani et al., 2017), whose self-attention mechanism affords bidirectional contextualization while maintaining high parallelizability. However, the quadratic time and memory costs of full self-attention remain a central bottleneck (Tay et al., 2022; Munkhdalai et al., 2024), limiting practical extension of context windows in cost-sensitive deployments. A large body of work has sought to mitigate this bottleneck (e.g., via using linear attention (Katharopoulos et al., 2020; Choromanski et al., 2021; Peng et al., 2025) and RNN-inspired architectures (Gu et al., 2021; Gupta et al., 2022; Fu et al., 2022; Gu & Dao, 2023)), but little of it has been adapted to the bidirectional, encoder-only paradigm. Meanwhile, BERT itself was modernized through larger pretraining corpora, architectural refinements (e.g., FlashAttention (Dao et al., 2022), SwiGLU activations (Shazeer, 2020), and RoPE positional encoding (Su et al., 2021)), and new pretraining and fine-tuning strategies (Liu et al., 2019; Portes et al., 2023; Warner et al., 2025), among others.

Most recently, Avey (Hammoud & Acharya, 2025) was introduced as an autoregressive architecture that departs from both Transformer- and RNN-based designs. It partitions a sequence into splits, ranks and retrieves the most relevant splits for each target split, and applies a dynamically parameterized neural processor to contextualize them. By decoupling context width from sequence length, Avey enables efficient long-range interactions and extrapolation far beyond its training window, facilitating practical context extensions at realistic compute budgets.

To elaborate on its technical aspects, Avey is a recurrence- and attention-free architecture comprising two principal components, a ranker and a neural processor. The ranker slices each input sequence into splits of consecutive tokens and selects the top k most relevant splits for each current split being processed by the neural processor. The neural processor consists of three core units, the enricher, contextualizer, and fuser. The enricher enhances the quality of token embeddings by expanding their learnable features using a position-wise neural network. The contextualizer is an embedding-wise neural network with dynamic parameterization and cosine-similarity-based selectivity, enabling interactions between relevant tokens across the current and top k splits. Lastly, the fuser learns a function that integrates the contextualized features produced by the contextualizer with some uncontextualized features bypassed by a partial-embedding bypassing mechanism.

Although originally formulated for causal language modeling, Avey’s cosine-similarity-based selectivity and learned cross-embedding linear transformation make it naturally amenable to a bidirectional, encoder-style adaptation. In this paper, we introduce **Avey-B**, a bidirectional reformulation of Avey for the encoder-only setting, and compare it empirically against widely used and recently introduced Transformer-based encoders, namely, BERT (Devlin et al., 2019), RoBERTa (Liu et al., 2019), ModernBERT (Warner et al., 2025), and NeoBERT (Breton et al., 2025). We further propose architectural advances that enhance its effectiveness and efficiency, including: (1) decoupled static and dynamic parameterizations; (2) a stability-oriented, row-normalized similarity scores in the dynamic layers; and (3) a compression module that reduces retrieved context before contextualization in the neural processor.

Concretely, Avey-B contextualizes tokens via either a learned static linear projection *or* a dynamic similarity matrix computed from cosine similarities, *in any given layer*. This contrasts with Avey, which multiplicatively couples the learned projection with cosine scores element-wise *in every layer*. By decoupling the static and dynamic parameterizations, Avey-B avoids destructive interactions between fixed weights and input-dependent scores, most notably inversion effects where a token highly similar to a neuron’s current token is forced to contribute less than a less-similar one. In addition, we normalize the cosine scores at each position by the sum of that position’s scores over all tokens, stabilizing training and consistently improving downstream task performance.

Alongside, we observe that extending Avey to the bidirectional paradigm without some modifications may introduce a scalability issue. Specifically, in the original design of Avey, each split is concatenated with its top- k relevant splits and jointly contextualized in a single pass through the neural processor. Performing this for *every* split inflates the input size to roughly k times the number of tokens, substantially increasing processing time. In the autoregressive regime, this overhead is mitigated by training on short context widths, leveraging Avey’s ability to extrapolate well beyond that. It is also tolerable at inference because *only* the most recent split is contextualized with its top- k splits to generate the next token. In the bidirectional inference setting, however, this strategy is infeasible, since *all* splits must be contextualized to produce complete token-level representations.

Building on this observation and recognizing that inference efficiency is critical for encoder models (especially in industry where they are commonly used (Raghavan, 2020; Zhu, 2019; Guo et al., 2020; Warner et al., 2025)), we introduce a neural compression scheme in the ranker. More precisely, we compress each split together with its top- k retrieved splits back to the size of a single split via a learned linear projection. As a result, the neural processor contextualizes only as many tokens as in the original input sequence, avoiding redundant computations over the appended top- k splits. Because the neural processor operates on each split independently, Avey-B achieves higher throughput than Transformer-based encoders, while preserving high accuracy across a wide range of downstream benchmarks.

To summarize, our main contributions in this paper are as follows:

- We propose Avey-B, a bidirectional encoder architecture that capitalizes on Avey by decoupling static and dynamic parameterizations and introducing a lightweight normalization scheme for dynamic contextualization.
- We redesign Avey’s ranker to compress each split’s top- k retrieved context into a fixed token budget, making the neural processor’s per-split compute independent of k while preserving the benefits of retrieving larger relevant token sets via increasing k .

- 108 • We conduct extensive design-choice and ablation studies to identify the most effective ar-
109 chitectural configuration and demonstrate how each proposed idea contributes to the per-
110 formance gains of Avey-B over the original Avey architecture.
- 111 • We show that Avey-B outperforms BERT (Devlin et al., 2019) and NeoBERT (Bret-
112 ton et al., 2025) across all the evaluated benchmarks, and consistently surpasses
113 RoBERTa (Liu et al., 2019) and ModernBERT (Warner et al., 2025) on token-classification
114 and information-retrieval tasks, despite being pretrained on $\sim 11 \times$ fewer tokens than Mod-
115 ernBERT.
- 116 • We illustrate that Avey-B scales efficiently with sequence length, yielding substantially
117 lower latency than Transformer-based encoders. Across 128–96 K tokens, Avey-B is con-
118 sistantly faster than all the evaluated Transformer baselines, and its advantage *widens* with
119 sequence length N . For example, at $N = 96$ K, Avey-B outpaces ModernBERT and
120 NeoBERT by $3.38 \times$ and $11.63 \times$, respectively.
- 121 • A power-law fit, $T(N) \propto N^{-\alpha}$, yields a markedly smaller decay exponent for Avey-B
122 ($\alpha \approx 0.44$) compared to ModernBERT ($\alpha \approx 0.77$) and NeoBERT ($\alpha \approx 0.81$), indicating
123 that Avey-B’s throughput decreases at only about half the rate of ModernBERT (and even
124 more slowly relative to NeoBERT) as sequence length increases.
- 125 • We release the full implementation and pretrained checkpoints of Avey-B (see Section 7),
126 enabling reproducibility and fostering future research.

128 The rest of the paper is organized as follows. Section 2 reviews related work, and Section 3 provides
129 background on Avey. We elaborate on the Avey-B architecture, including its bidirectional contextu-
130 alization, decoupled parameterization, and neural compression in Section 4. Section 5 presents our
131 experimental setup, design choices and ablations, and effectiveness and efficiency results. Finally,
132 we conclude in Section 6.

133 2 RELATED WORK

134 The introduction of GPT (Radford et al., 2018) in 2018 marked a turning point in large-scale lan-
135 guage modeling, establishing the now-standard paradigm of pretraining Transformer-based models
136 on massive unlabeled corpora followed by supervised fine-tuning on task-specific data. GPT op-
137 timized a causal language modeling (CLM) objective, pretraining a unidirectional, decoder-only
138 Transformer (Vaswani et al., 2017) for next-token prediction. The resulting pretrained model can
139 then be effectively fine-tuned with modest labeled data to a broad range of downstream tasks, in-
140 cluding text classification (Wang et al., 2019b), natural language inference (Bowman et al., 2015;
141 Williams et al., 2018), and question answering (Rajpurkar et al., 2016; Lai et al., 2017), to mention
142 just a few. This pretrain–fine-tune paradigm yielded state-of-the-art performance on these tasks at
143 the time (Radford et al., 2018).

144 BERT (Devlin et al., 2019) extended this paradigm by replacing the unidirectional decoder with a
145 fully bidirectional encoder. Concretely, it introduced two pretraining objectives, masked language
146 modeling (MLM), which reconstructs randomly masked tokens in an input sequence, and next sen-
147 tence prediction (NSP), which models inter-sentence relationships. By contextualizing tokens in
148 both directions, BERT delivered substantial gains over causally pretrained models, particularly on
149 benchmarks such as GLUE (Wang et al., 2019b), MultiNLI (Williams et al., 2018), and SQuAD (Ra-
150 jpurkar et al., 2016), among others.

151 RoBERTa (Liu et al., 2019) robustly optimized BERT by retaining its overall architecture while
152 systematically revisiting nearly every aspect of its pretraining setup. Key modifications included
153 removing the NSP objective, pretraining with larger batches and longer sequences, adopting dy-
154 namic masking strategies, and scaling to substantially larger corpora. Building on this foundation,
155 DeBERTa (He et al., 2021b;a; 2023) introduced disentangled attention, which separates content and
156 positional information into distinct attention matrices, and improved fine-tuning stability through
157 virtual adversarial training. Together, these innovations further advanced performance on some
158 challenging benchmarks such as SuperGLUE (Wang et al., 2019a).

159 Subsequent work emphasized both architectural refinements and pretraining efficiency. For ex-
160 ample, MosaicBERT (Portes et al., 2023) integrated FlashAttention (Dao et al., 2022), ALiBi positional
161 biases (Press et al., 2022), and gated linear units (GLU) (Dauphin et al., 2017; Shazeer, 2020) to ac-
celerate pretraining while maintaining strong downstream accuracy. NomicBERT (Nussbaum et al.,

2024) adopted SwiGLU (Shazeer, 2020) and rotary positional encodings (RoPE) (Su et al., 2021).
 163 NeoBERT (Breton et al., 2025) combined RoPE, SwiGLU, and RMSNorm (Zhang & Sennrich,
 164 2019) with depth-width rebalancing and large-scale pretraining. ModernBERT (Warner et al., 2025)
 165 pushed this trend further, employing many of these techniques (e.g., RoPE, FlashAttention, and al-
 166 ternating global/local attention), supporting context windows for up to 8,192 tokens, and pretraining
 167 on multi-trillion-token corpora.

168 All of the above models are Transformer-based, leveraging self-attention to provide effective bidi-
 169 rectional contextualization while maintaining high pretraining parallelism. Recently, a fundamen-
 170 tally different architecture named Avey (Hammoud & Acharya, 2025) was introduced. Avey is
 171 attention-free and can process virtually unlimited sequence lengths (see Section 3). Avey-B capital-
 172 izes on Avey to support bidirectional contextualization, mirroring the shift from GPT-style decoder-
 173 only to BERT-style encoder-only models in the Transformer family. We empirically compare Avey-
 174 B against BERT, RoBERTa, ModernBERT, and NeoBERT in Section 5.

175 3 BACKGROUND

177 The original Avey architecture decouples sequence length from context width by pairing a
 178 lightweight ranker with a data-dependent neural processor. We next provide a background on both.

179 3.1 RANKER

180 Avey partitions an input sequence of length N into equal-sized *splits* of S tokens, applying zero-
 181 padding if N is not divisible by S . For a given *current* split, the ranker computes its relevance to
 182 each preceding split using the MaxSim operator (Khattab & Zaharia, 2020), orders them by their
 183 MaxSim scores, and selects the top- k splits for contextualization.

184 Before contextualization, the MaxSim scores of the top- k selected splits are normalized by dividing
 185 each score by the maximum among them. Each selected split is then weighted by its normalized
 186 score and concatenated with the current split. This *weighted-selective-split* mechanism prunes irrel-
 187 evant global tokens and scales the contribution of each retrieved split based on relevance.

188 Crucially, the ranker is invoked only *once* per full forward/backward pass, independent of the num-
 189 ber of neural-processor layers. Matching each split against all preceding splits yields a training-time
 190 compute cost of $\mathcal{O}(N^2d)$, where d is the embedding dimension.

192 3.2 NEURAL PROCESSOR

193 The neural processor ingests the current split and its weighted top- k retrieved splits, and processes
 194 them through a layered architecture. Each layer comprises three modules, an **enricher**, a **contextu-**
 195 **alizer**, and a **fuser**.

197 The enricher is a single-layer, position-wise neural network applied independently to each token
 198 embedding. Given C input embeddings arranged as $\mathbf{X} \in \mathbb{R}^{C \times d}$, the enricher computes a matrix
 199 $\mathbf{Z} \in \mathbb{R}^{C \times m}$ (with $m > d$) as follows:

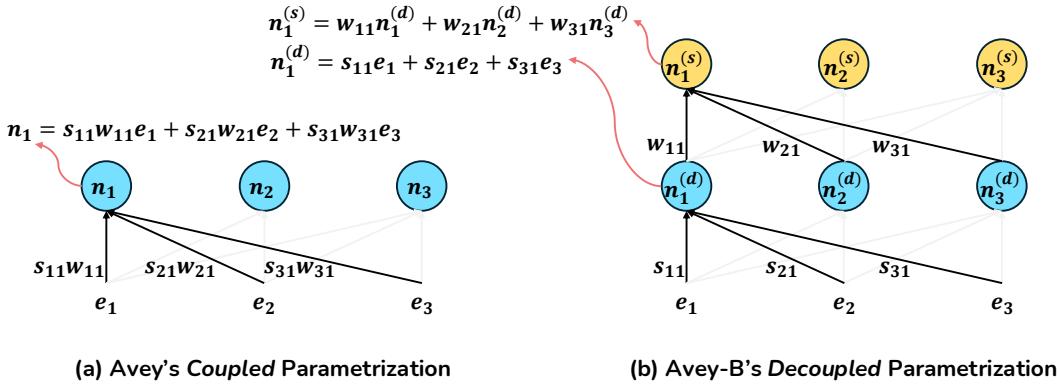
$$200 \quad \mathbf{Z} = \sigma(\mathbf{X}\mathbf{U} + \mathbf{b}), \quad (1)$$

202 where $\mathbf{U} \in \mathbb{R}^{d \times m}$ is a learnable weight matrix, $\mathbf{b} \in \mathbb{R}^{C \times m}$ represents biases, and $\sigma(\cdot)$ is an
 203 activation function. The output \mathbf{Z} is partitioned into a *head* $\mathbf{Z}_h \in \mathbb{R}^{C \times m_h}$, which is bypassed
 204 directly to the fuser, and a *tail* $\mathbf{Z}_t \in \mathbb{R}^{C \times m_t}$ (with $m = m_h + m_t$), which is forwarded to the
 205 contextualizer. This *partial-embedding bypassing* technique preserves raw token-specific features
 206 and mitigates degradation effects (e.g., over-smoothing), as the number of layers is increased.

207 The contextualizer operates on the tail \mathbf{Z}_t . Each m_t -dimensional tail embedding is split evenly into a
 208 *gating* left half and a *contextual* right half, yielding $\mathbf{Z}_{tl} \in \mathbb{R}^{C \times \frac{m_t}{2}}$ and $\mathbf{Z}_{tr} \in \mathbb{R}^{C \times \frac{m_t}{2}}$, respectively.
 209 Formally, the contextualizer is a single-layer, embedding-wise network that updates \mathbf{Z}_{tr} as follows:

$$211 \quad \mathbf{c}(\mathbf{Z}_t) = \mathbf{Z}_{tl} \odot \sigma\left((\mathbf{V} \odot \mathcal{N}(\mathbf{Z}_{tr}) \mathcal{N}(\mathbf{Z}_{tr})^\top) \mathbf{Z}_{tr} + \mathbf{b}'\right), \quad (2)$$

214 where $\mathbf{V} \in \mathbb{R}^{C \times C}$ is a learned cross-embedding matrix, \odot denotes element-wise (Hadamard) mul-
 215 tiplication, $\mathcal{N}(\cdot)$ applies row-wise ℓ_2 normalization (so $\mathcal{N}(\mathbf{Z}_{tr}) \mathcal{N}(\mathbf{Z}_{tr})^\top$ computes cosine similar-
 216 ties between embeddings), and \mathbf{b}' is an optional bias. Intuitively, each neuron aggregates statically



229
230
231
232
233
234
235
236
237
238
239
240
241
242
243
244
245
246
247
248
249
250
251
252
253
254
255
256
257
258
259

Figure 1: A simple illustration of coupled (a) and decoupled (b) parameterizations (e_i = embedding i ; s_{ij} = cosine similarity score between e_i and e_j ; n_i = neuron i , $n_i^{(d)}$ = neuron i in dynamic layer d ; $n_i^{(s)}$ = neuron i in static layer s ; and w_{ij} = weight corresponding to e_i or $n_i^{(d)}$ used in the weighted sum of n_j or $n_j^{(s)}$, respectively).

and dynamically weighted contributions from other embeddings, and the resulting update is gated by \mathbf{Z}_{tl} . The learned matrix \mathbf{V} provides position-sensitive mixing, so no additional positional encodings are required within the contextualizer.

The fuser combines the bypassed and contextualized streams and projects the output back to the model embedding dimension d as follows:

$$f(\mathbf{Z}) = [\mathbf{Z}_h \parallel \mathbf{c}(\mathbf{Z}_t)] \mathbf{O}, \quad (3)$$

where $\mathbf{O} \in \mathbb{R}^{(m_h+m_t/2) \times d}$ is a learned projection matrix. As with the enricher, the fuser is applied independently to each token embedding. Its output is merged with the enricher's input within the same layer via a residual, element-wise addition.

Aggregating the costs of the ranker, enricher, contextualizer, and fuser over L layers yields a training complexity of $\mathcal{O}(N^2d)$ for an input sequence of length N . At inference, only the most recent split is contextualized for autoregressive decoding, reducing the complexity to $\mathcal{O}(N)$.

4 AVEY-B

Avey-B is a bidirectional reformulation of Avey. We next elaborate on its architecture and computational implications.

4.1 BIDIRECTIONAL CONTEXTUALIZATION

Avey-B drops the autoregressive mask in Avey's contextualizer, allowing each token representation to condition on both left and right contexts. Specifically, when a split is contextualized with its top- k selected splits, all token interactions are permitted, without any causal masking. This converts Avey into an encoder-style architecture while preserving selective global access via the ranker.

4.2 DECOUPLED PARAMETRIZATION

In Avey, the contextualizer multiplies (element-wise) a learned, static weight matrix (i.e., \mathbf{V}) with an input-dependent cosine-similarity matrix (i.e., the outcome of $\mathcal{N}(\mathbf{Z}_{tr})\mathcal{N}(\mathbf{Z}_{tr})^\top$), then uses the result to linearly combine the input embeddings (i.e., \mathbf{Z}_{tr}) into contextualized representations (see Equation 2). This tight coupling of fixed parameters with data-driven relevance scores can induce pathological behaviors. For instance, a token that is highly similar to a neuron's current token (thus, yielding a positive cosine score) can be forced to contribute negatively whenever the corresponding entry of \mathbf{V} is negative. Consequently, the neuron's update violates *monotonicity with respect to relevance*, which entails that a more-relevant token must contribute at least as much as a less-relevant one, and increasing a token's relevance must not reduce or invert its contribution.

Fig. 1 (a) offers a simple illustration of the problem. Suppose the cosine similarity s_{21} exceeds s_{31} ; then embedding e_2 should contribute at least as much to neuron n_1 's update as embedding e_3 .

270 In Avey’s coupled design, the element-wise product with learned weights can invert this ordering.
 271 Specifically, if weight $w_{31} \gg$ weight w_{21} (or the signs differ), the effective contributions $w_{21}s_{21}$
 272 and $w_{31}s_{31}$ can be reversed (especially at inference), undermining evidence accumulation from the
 273 most informative tokens.

274 Avey-B addresses this problem by *decoupling* the two parameter sources (i.e., learned weights and
 275 input-driven similarities) and interleaving them across depth (we compare different interleaving pat-
 276 terns in Appendix D). In particular, we define each layer to be either *static* or *dynamic*. A static layer
 277 applies a learned linear transformation on embeddings while a dynamic layer weights them solely by
 278 cosine similarity. Alternating these layers preserves monotonicity of similarity-based updates with
 279 respect to relevance (i.e., when $s_{21} > s_{31}$, token 2 receives no smaller, and typically larger, positive
 280 influence than token 3 in the dynamic layer).

281 Fig. 1 (b) demonstrates the effect of decoupled parametrization. If $s_{21} > s_{31}$, the dynamic layer
 282 for neuron $n_1^{(d)}$ assigns a larger normalized weight to e_2 than to e_3 , with no learned weights in-
 283 tervening in this similarity-based update. In the following static layer (assuming an interleaved
 284 dynamic-static pattern) for $n_1^{(s)}$, both contributions are scaled by the same coefficient w_{11} ; hence,
 285 the ordering established by the dynamic layer is preserved. More generally, because the static layer
 286 is similarity-agnostic, it cannot retroactively modify the normalized dynamic scores or introduce
 287 similarity-conditioned sign flips.

288 To this end, decoupling static and dynamic computations maintains the *monotonicity guarantee* at
 289 each dynamic layer while still allowing representation shaping in static layers. Static layers will, of
 290 course, change the representations from which subsequent similarities are computed, but they do not
 291 alter the scores already assigned by a preceding dynamic layer and thus do not violate monotonicity
 292 for that layer. We provide a formal proof for this monotonicity guarantee in Appendix A and ablate
 293 this decoupling design choice in Appendix H. We further analyze the coupled versus decoupled
 294 parameterizations and their implications in Appendix L.

295 Formally, let $\mathbf{Z}_{\text{tr}} \in \mathbb{R}^{C \times d'}$ denote the matrix of contextual, right-tail components for C enriched
 296 embeddings, where d' is the contextualizer’s right-tail dimension (see Section 3 for more information
 297 on all notations). In Avey-B, a static layer applies a learned cross-embedding linear transformation
 298 as follows:

$$300 \quad \mathbf{c}_{\text{static}}(\mathbf{Z}) = \sigma(\mathbf{V} \mathbf{Z}_{\text{tr}} + \mathbf{b}^{(s)}), \quad (4)$$

303 where $\mathbf{V} \in \mathbb{R}^{C \times C}$ is a learned matrix, $\mathbf{b}^{(s)} \in \mathbb{R}^{C \times d'}$ is an optional bias, and $\sigma(\cdot)$ is an activation
 304 function. Intuitively, each neuron first aggregates linearly the C embeddings and then applies the
 305 pointwise activation σ .

307 On the flip side, an Avey-B’s dynamic layer computes an input-dependent similarity matrix from
 308 \mathbf{Z}_{tr} and utilizes it to mix embeddings as follows:

$$310 \quad \mathbf{S} = \mathcal{N}(\mathbf{Z}_{\text{tr}}) \mathcal{N}(\mathbf{Z}_{\text{tr}})^{\top} \in \mathbb{R}^{C \times C}, \quad (5)$$

$$312 \quad \tilde{\mathbf{S}}_{i,j} = \frac{\mathbf{S}_{i,j}}{\sum_{j=1}^C \mathbf{S}_{i,j} + \varepsilon} \quad (\text{row-wise sum normalization}), \quad (6)$$

$$315 \quad \mathbf{c}_{\text{dyn}}(\mathbf{Z}) = \sigma(\tilde{\mathbf{S}} \mathbf{Z}_{\text{tr}} + \mathbf{b}^{(d)}). \quad (7)$$

317 Here $\mathcal{N}(\cdot)$ denotes per-row ℓ_2 normalization to unit length so that \mathbf{S} encodes cosine similarities;
 318 $\varepsilon > 0$ is a small stabilizer ensuring a positive denominator; $\mathbf{b}^{(d)} \in \mathbb{R}^{C \times d'}$ is an optional bias, and $\tilde{\mathbf{S}}$
 319 is a simple *sum-normalized* similarity matrix. This row-wise normalization yields a row-stochastic
 320 similarity operator (row sums ≤ 1), which bounds per-row gain and mitigates the growth of large
 321 singular values through depth, improving numerical conditioning and trainability. In the unnor-
 322 malized case (i.e., in Avey), inflated singular values can drive activation and gradient growth with
 323 depth, resulting in unstable optimization and degraded generalization. We ablate this normalization
 technique and show consistent gains over softmax-based and RMS-style alternatives in Appendix E.

324 4.3 NEURAL COMPRESSION
325

326 A key bottleneck in extending Avey to the bidirectional setting is the per-split concatenation strategy,
327 whereby each current split is concatenated with its top- k retrieved splits for contextualization. This
328 is manageable in Avey’s autoregressive inference because only the most recent split is expanded. In
329 the bidirectional regime, however, every split within the contextualizer’s window must be expanded,
330 which inflates the effective sequence length by a factor of $k + 1$.

331 To mitigate this, Avey-B introduces a neural compressor within the ranker to condense the con-
332 catenated $(k + 1)S$ -token block back to S tokens, where S denotes the split size. Specifically, the
333 compressor is an embedding-wise neural module that maps the $(k + 1)S$ input tokens to S represen-
334 tative tokens, effectively distilling cross-split information before it is passed to the neural processor.
335 To preserve signal from the block’s current split, Avey-B adds a residual connection between the
336 compressor output and the split’s original S tokens, which improves stability and downstream ef-
337 fectiveness. An ablation studying the impact of this residual on Avey-B’s accuracy is presented in
338 Appendix H.

339 Formally, let $\mathbf{X}_{\text{cat}} \in \mathbb{R}^{(k+1)S \times d}$ be the concatenation of a *single* split with its top- k retrieved splits,
340 where d denotes the embedding dimension. Subsequently, Avey-B produces a compressed output
341 $\hat{\mathbf{X}} \in \mathbb{R}^{S \times d}$ as follows:

$$342 \quad \hat{\mathbf{X}} = \mathbf{P} \mathbf{X}_{\text{cat}} \quad (8)$$

343 where $\mathbf{P} \in \mathbb{R}^{S \times (k+1)S}$ is a learnable matrix that performs a linear cross-token transformation, and $\hat{\mathbf{X}}$
344 replaces \mathbf{X}_{cat} as the input to the neural processor. Because \mathbf{P} is a learned matrix, the compressor can
345 preserve globally informative content while discarding potential redundancy, yielding a favorable
346 accuracy/throughput trade-off. We study the effect of the compressor on Avey-B’s accuracy and
347 throughput in Appendix H.

348 As discussed in (Hammoud & Acharya, 2025), computation in the neural processor largely domi-
349 nates that of the ranker. As such, when Avey-B reduces the number of tokens contextualized per
350 split from $(k+1)S$ to S , throughput improves by $4.37\times$ (see Fig. 3 in Appendix H), albeit leaving
351 Avey’s asymptotic complexity unchanged (still quadratic with respect to the sequence length N).
352

353 5 EXPERIMENTS

354 5.1 EXPERIMENTAL SETUP

355 In this section, we compare Avey-B with widely used and recently introduced Transformer-based
356 encoders, namely, BERT (Devlin et al., 2019), RoBERTa (Liu et al., 2019), ModernBERT (Warner
357 et al., 2025), and NeoBERT (Breton et al., 2025). We evaluate two Avey-B model sizes, *base*
358 and *large*, each pretrained on 180B tokens drawn from the FineWeb corpus (Hugging Face, 2023).
359 Pretraining details and information about all the evaluated models are provided in Appendix B.

360 To assess effectiveness, we adopt the evaluation protocol of Boukhleff *et al.* (Gisserot-Boukhleff
361 et al., 2025), targeting breadth across four downstream categories prevalent in practice, includ-
362 ing Sequence Classification (SC), Token Classification (TC), Question Answering (QA), and In-
363 formation Retrieval (IR). Each category is represented by three established benchmarks, namely,
364 MNLI (Williams et al., 2018), QQP (Wang et al., 2017), and SST-2 (Socher et al., 2013) under SC;
365 CoNLL-2003 (Sang & De Meulder, 2003), OntoNotes (Hovy et al., 2006), and UNER (Mayhew
366 et al., 2023) under TC; ReCoRD (Wang et al., 2019a), SQuAD (Rajpurkar et al., 2016), and SQuAD-
367 v2 (Rajpurkar et al., 2018) under QA; and MLDR (Multi-Granularity, 2024), MS MARCO (Bajaj
368 et al., 2016), and NQ (Kwiatkowski et al., 2019) under IR.

369 We fine-tuned benchmarks under SC and TC for 1 epoch, QA for 4 epochs, and IR for 1,000 op-
370 timization steps. For each benchmark, we swept four learning rates $\{2 \times 10^{-5}, 6 \times 10^{-5}, 1 \times$
371 $10^{-4}, 5 \times 10^{-4}\}$ and trained each configuration with 10 independent random seeds. Akin to (Liu
372 et al., 2019), the reported results for each model are the *median* scores across seeds at the *best*
373 learning rate. SC is evaluated with accuracy, TC and QA with F1 score, and IR with NDCG@10. Lastly,
374 following Boukhleff *et al.* (Gisserot-Boukhleff et al., 2025), we used linear learning-rate decay with
375 warmup over the first 10% of steps.

378 Table 1: Design and masked language modeling (MLM) choices.
379

380 Question	381 Answer	382 Experiments
383 What is the most effective arrangement of static (S) and 384 dynamic (D) layers?	385 Interleaved layers with a 386 repeating $S \rightarrow D$ pattern	387 Appendix D
388 What is the most effective normalization technique 389 within the dynamic layers, Softmax, RMS Norm, or 390 Row-wise normalization?	391 Row-wise normalization	392 Appendix E
393 What are the best values for sequence length N , split 394 size S , and top k splits?	395 $N = 2048, S = 256, k = 3$	396 Appendix F
397 Should the ranker operate bidirectionally as well?	398 No	399 Appendix C
400 What is the best masking rate?	401 20%	402 Appendix G

403 Table 2: Effectiveness results for several encoders at different scales (M = Medium).
404

405 Model	406 SC				407 TC				408 QA				409 IR				
	MNLI	QQP	SST-2	Avg.	CONLL	Onto.	UNER	Avg.	ReCoRD	SQuAD	SQuAD	Avg.	MLDR	MS	NQ	Avg.	
										v2				MARCO			
410 Base	Avey-B	83.58	89.81	92.94	88.78	92.88	93.80	94.10	93.59	44.03	74.44	68.88	62.45	63.83	88.14	83.62	78.53
	BERT	81.92	88.57	90.94	87.14	90.25	91.03	88.20	89.82	36.76	72.20	63.99	57.65	57.42	81.15	80.66	73.08
	RoBERTa	86.42	89.12	92.78	89.44	90.55	92.11	88.16	90.27	67.86	80.68	76.62	75.05	56.07	86.47	80.30	74.28
	ModernBERT	86.72	89.81	92.32	89.61	92.30	93.74	92.30	92.78	65.73	80.23	77.36	74.44	54.29	88.09	75.24	72.54
411 M	NeoBERT	82.53	88.88	84.69	85.36	87.55	88.88	88.17	88.20	37.74	64.84	64.42	55.67	39.98	70.76	59.43	56.72
	Avey-B	85.66	89.22	94.38	89.75	93.60	94.09	94.32	94.00	58.22	77.30	72.46	69.32	67.05	88.72	86.24	80.67
	BERT	85.08	89.27	92.26	88.87	88.54	90.71	86.09	88.44	52.02	77.93	72.96	67.64	61.08	87.71	85.42	78.07
	RoBERTa	90.16	89.49	94.67	91.44	91.71	92.70	88.79	91.07	80.86	84.00	83.04	82.63	58.50	89.43	85.91	77.95
412 Large	ModernBERT	90.53	90.73	95.99	92.41	92.43	93.79	92.92	93.05	73.05	82.02	79.96	78.34	59.64	88.82	81.36	76.61

413 To measure efficiency, we report *latency* (seconds/forward pass) and *throughput* (tokens/second) as
414 a function of input context length (with a fixed batch size of 8), benchmarking models on NVIDIA
415 H200 GPUs under identical software and precision settings. These measurements provide a
416 direct comparison of scalability and deployment efficiency between Avey-B and Transformer-based
417 encoders.

418 5.2 DESIGN CHOICES AND ABLATIONS

419 Following the methodology established in the original Avey work (Hammoud & Acharya, 2025), we
420 conduct systematic design-choice studies to identify effective configurations for Avey-B. We also
421 sweep the masking rate to select a robust setting for pretraining Avey-B. Table 1 summarizes these
422 studies and links to the supporting experiments for each conclusion.

423 Moreover, we perform ablation studies to quantify the contribution of each refinement to Avey-B.
424 Specifically, we measure incremental gains from: (1) decoupling and interleaving static and dynamic
425 parameterizations, (2) introducing normalization within the dynamic layers, and (3) integrating a
426 neural compressor into the ranker. Appendix H presents the full results and analysis.

427 5.3 EFFECTIVENESS

428 We now evaluate Avey-B on standard SC, TC, QA, and IR benchmarks, comparing it against BERT,
429 RoBERTa, ModernBERT, and NeoBERT. For all baselines we consider *base* and *large* configurations
430 (see Table 3 in Appendix B), except for NeoBERT which is evaluated on its only publicly
431 available size, that is, *medium*. Table 2 summarizes all the results.

432 At the *base* scale (and *medium* for NeoBERT), Avey-B surpasses BERT and NeoBERT across all
433 task categories, despite using $\sim 85M$ fewer parameters than NeoBERT. It also delivers the strongest
434 results on TC and IR, outperforming *all* Transformer-based models in both categories. For SC,
435 Avey-B attains the best scores on QQP (tied with ModernBERT) and SST-2, while trailing RoBERTa
436 and ModernBERT on MNLI. For QA, Avey-B leads on SQuAD-v2 but lags RoBERTa and Modern-
437 BERT on ReCoRD and SQuAD.

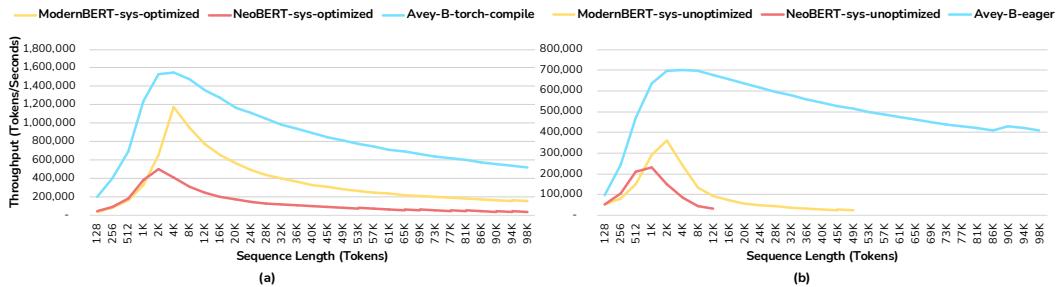


Figure 2: Throughput of Avey-B, ModernBERT, and NeoBERT on NVIDIA B200 GPUs with mixed precision (BF16). We use Avey-B *base*, ModernBERT *base*, and NeoBERT *medium* (the only publicly available size). Avey-B is shown in (a) as optimized using `torch.compile` (no fused-kernel implementation is available yet) and in (b) as unoptimized (eager). For ModernBERT and NeoBERT, throughput is shown for system-optimized (with FlashAttention) and system-unoptimized (eager) variants in (a) and (b), respectively.

At the *large* scale (and *medium* for NeoBERT), Avey-B again outperforms BERT and NeoBERT across all task categories. It also offers the strongest results on TC and IR, surpassing *every* Transformer-based model. Notably, the Avey-B *base* model even exceeds all *large* Transformer encoders on TC and IR (despite also being pretrained on $\sim 11 \times$ fewer tokens than ModernBERT, for example). These results highlight Avey-B’s advantage for both local, span-sensitive decisions (TC) and long-document encoding (IR).

In summary, at both *base* and *large* scales, Avey-B exceeds BERT and NeoBERT on every evaluated benchmark and delivers consistent gains over all baselines on TC and IR. We attribute these gains to two design factors: (1) TC tasks hinge on local evidence within short spans, and Avey-B’s split-based processing with pruning of low-relevance splits/tokens sharpens the signal-to-noise ratio; and (2) IR benefits from selectively pairing globally relevant content with its immediate local context when encoding long documents (an inductive bias Avey-B enforces by construction), whereas full bidirectional processing over all tokens (like in all the Transformer-based models) tends to admit distractors and dilute relevance as sequence length grows.

5.4 EFFICIENCY

Avey is a recent architecture and still lacks a fused-kernel (CUDA/Triton) implementation. Consequently, we evaluate its inference efficiency using `torch.compile`, which performs graph capture and backend code generation but stops short of handcrafted, specialized fused kernels available to mature Transformer encoders. We denote this configuration as *Avey-B-torch-compile*. To quantify compilation gains and offer a reference without compiler optimizations, we also report Avey-B’s efficiency in the *eager* PyTorch mode, referred to as *Avey-B-eager*. For Transformer baselines, we employ ModernBERT and NeoBERT as representative encoders, especially since they were both recently modernized and optimized using FlashAttention (Dao et al., 2022), RoPE positional encoding (Su et al., 2021), and several other engineering techniques (see Section 2).

RoPE enables evaluation beyond a model’s pre-trained context window, which we leverage to measure throughput and latency for ModernBERT and NeoBERT at long sequence lengths. Although such extrapolation may degrade task effectiveness, this does not impact our study, as we are concerned solely with efficiency rather than model quality. In contrast, Avey does not face this limitation because its architecture decouples context width from sequence length, allowing Avey-B to operate at arbitrarily long sequences *without* any additional pre-training. Accordingly, we extend sequence length as needed and compare Avey-B *base* against ModernBERT *base* and NeoBERT *medium* (the only publicly available variant), reporting throughput in this section and latency in Appendix I. The qualitative conclusions presented in this section, however, apply to both throughput and latency.

For ModernBERT and NeoBERT, we also report throughput under both optimized and unoptimized settings, corresponding to execution *with* and *without* FlashAttention. We refer to these configurations as *sys-optimized* and *sys-unoptimized*, respectively. As shown in Fig. 2, the throughput curves for all encoders across sequence lengths from 128 to 96k tokens and under both settings exhibit a consistent structure. In particular, throughput increases at short sequences, plateaus at

486 medium lengths, and eventually declines as sequence lengths grow large. This behavior follows
 487 from the fixed batch size of eight used in our study. To elaborate, each forward pass incurs a con-
 488 stant memory-loading overhead, which dominates when few tokens are processed, leading to low
 489 GPU utilization. As sequence length increases, more tokens are processed per batch, amortizing
 490 this communication cost and improving utilization. At sufficiently large sequence lengths, how-
 491 ever, computation becomes the dominant bottleneck due to arithmetic intensity, memory-bandwidth
 492 saturation, and large intermediate activations, causing throughput to decline.

493 Despite sharing this qualitative profile, the encoders differ markedly in their quantitative scaling.
 494 We characterize long-context throughput using a power-law decay model, $T(N) \propto N^{-\alpha}$, where N
 495 denotes the sequence length and smaller exponents α indicate better long-context efficiency. Un-
 496 der optimized settings (see Fig. 2 (a)), ModernBERT-sys-optimized and NeoBERT-sys-optimized
 497 demonstrate decay exponents of $\alpha_{\text{ModernBERT}} = 0.77$ and $\alpha_{\text{NeoBERT}} = 0.81$, respectively, consistent
 498 with the bandwidth- and memory-driven limitations inherent to quadratic self-attention. In con-
 499 trast, Avey-B-torch-compile displays a substantially milder decay with exponent $\alpha_{\text{Avey-B}} = 0.44$,
 500 sustaining significantly higher throughput across the long-context regime.

501 The unoptimized measurements further accentuate these differences (see Fig. 2 (b)).
 502 ModernBERT-sys-unoptimized and NeoBERT-sys-unoptimized exhibit considerably steeper decay,
 503 with exponents $\alpha_{\text{ModernBERT}} = 1.03$ and $\alpha_{\text{NeoBERT}} = 1.30$, and both encounter out-of-memory fail-
 504 ures well before the maximum tested sequence length. Conversely, Avey-B-eager again achieves
 505 the mildest decay, with $\alpha_{\text{Avey-B}} = 0.33$, and maintains stable throughput across the entire sequence-
 506 length range. These results indicate that Avey-B’s scaling advantage is structural rather than an
 507 artifact of kernel-level or compiler optimizations. Specifically, its neural processor is inherently
 508 less sensitive to sequence length because computation depends on the fixed split size S , not the full
 509 sequence length N . With N/S splits, the total processing cost is $(N/S) \times S^2 = NS = \mathcal{O}(N)$,
 510 yielding linear scaling in N and a growing throughput advantage at long contexts, while still being
 511 the fastest tested encoder at short contexts.

512 6 CONCLUSION

513 In this paper, we presented Avey-B, a bidirectional encoder built on Avey, a new attention-free
 514 foundation model. Avey-B contributes three architectural innovations: (1) decoupling static and
 515 dynamic parameterizations, (2) row-normalized similarity in the dynamic layers, and (2) a neural
 516 compression module for improving effectiveness and efficiency. Results show that Avey-B deliv-
 517 ers consistent gains over Transformer-based encoders, including BERT, RoBERTa, ModernBERT,
 518 and NeoBERT on token classification and information retrieval, while outperforming BERT and
 519 NeoBERT on every evaluated benchmark. These findings entail that attention might not be the only
 520 viable route to strong bidirectional encoders and motivate further study of retrieval-conditioned,
 521 non-attention architectures.

523 7 REPRODUCIBILITY

525 All results reported in this paper are reproducible. Section 4 specifies Avey-B’s components in detail.
 526 The full experimental methodology is provided in Section 5.1 and Appendix B. We attach a reposi-
 527 tory with code as supplementary material. The repository includes: (1) training, fine-tuning, and
 528 evaluation scripts; (2) configuration files with the exact hyperparameters used for every experiment;
 529 (3) data preprocessing instructions and dataset references/splits; and (4) environment specifications
 530 and run scripts to regenerate all tables and figures. Using the provided commands on hardware
 531 comparable to our setup reproduces the reported numbers within expected seed variance.

533 REFERENCES

536 answerdotai-base. Modernbert-base. <https://huggingface.co/answerdotai/>
 537 ModernBERT-base. Hugging Face model repository.

538 answerdotai-large. Modernbert-large. <https://huggingface.co/answerdotai/>
 539 ModernBERT-large. Hugging Face model repository.

540 Payal Bajaj, Daniel Campos, Nick Craswell, Li Deng, Jianfeng Gao, Xiaodong Liu, Rangan Ma-
 541 jumder, Andrew McNamara, Bhaskar Mitra, Tri Nguyen, et al. Ms marco: A human generated
 542 machine reading comprehension dataset. *arXiv preprint arXiv:1611.09268*, 2016.

543

544 Samuel R. Bowman, Gabor Angeli, Christopher Potts, and Christopher D. Manning. A large anno-
 545 tated corpus for learning natural language inference. In *Proceedings of the 2015 Conference on*
 546 *Empirical Methods in Natural Language Processing (EMNLP)*, 2015.

547 Lola Le Breton, Quentin Fournier, Mariam El Mezouar, John X. Morris, and Sarath Chandar.
 548 Neobert: A next-generation bert. *arXiv preprint arXiv:2502.19587*, 2025.

549

550 chandar-lab. Neobert. <https://huggingface.co/chandar-lab/NeoBERT>. Hugging
 551 Face model repository.

552 Krzysztof Choromanski, Valerii Likhoshesterov, David Dohan, Xingyou Song, Andreea Gane, Tamás
 553 Sarlós, Peter Hawkins, Jared Davis, Afroz Mohiuddin, Łukasz Kaiser, David Belanger, Lucy
 554 Colwell, and Adrian Weller. Rethinking attention with performers. In *International Conference*
 555 *on Learning Representations (ICLR)*, 2021.

556 Tri Dao, Daniel Y. Fu, Stefano Ermon, Atri Rudra, and Christopher Ré. Flashattention: Fast and
 557 memory-efficient exact attention with io-awareness. In *Advances in Neural Information Process-
 558 ing Systems (NeurIPS)*, 2022.

559

560 Yann N. Dauphin, Angela Fan, Michael Auli, and David Grangier. Language modeling with gated
 561 convolutional networks. In *Proceedings of the 34th International Conference on Machine Learn-
 562 ing (ICML)*, pp. 933–941. PMLR, 2017.

563 Jacob Devlin, Ming-Wei Chang, Kenton Lee, and Kristina Toutanova. Bert: Pre-training of deep
 564 bidirectional transformers for language understanding. *arXiv preprint arXiv:1810.04805*, 2019.

565

566 facebook-roberta-base. Roberta base (roberta-base). [https://huggingface.co/
 567 FacebookAI/roberta-base](https://huggingface.co/FacebookAI/roberta-base). Hugging Face model repository.

568

569 facebook-roberta-large. Roberta large (roberta-large). [https://huggingface.co/
 570 FacebookAI/roberta-large](https://huggingface.co/FacebookAI/roberta-large). Hugging Face model repository.

571

572 Daniel Y. Fu, Tri Dao, Khaled K. Saab, Armin W. Thomas, Atri Rudra, and Christopher Ré.
 573 Hungry hungry hippos: Towards language modeling with state space models. *arXiv preprint
 arXiv:2212.14052*, 2022.

574

575 Hippolyte Gisserot-Boukhlef, Nicolas Boizard, Manuel Faysse, Duarte M Alves, Emmanuel Mal-
 576 herbe, André FT Martins, Céline Hudelot, and Pierre Colombo. Should we still pretrain encoders
 577 with masked language modeling? *arXiv preprint arXiv:2507.00994*, 2025.

578

579 google-bert-base. Bert base uncased (bert-base-uncased). [https://huggingface.co/
 580 google-bert/bert-base-uncased](https://huggingface.co/google-bert/bert-base-uncased). Hugging Face model repository.

581

582 google-bert-large. Bert large uncased (bert-large-uncased). [https://huggingface.co/
 583 google-bert/bert-large-uncased](https://huggingface.co/google-bert/bert-large-uncased). Hugging Face model repository.

584

585 Albert Gu and Tri Dao. Mamba: Linear-time sequence modeling with selective state spaces. *arXiv
 586 preprint arXiv:2312.00752*, 2023.

587

588 Albert Gu, Karan Goel, and Christopher Ré. Efficiently modeling long sequences with structured
 589 state spaces. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2021.

590

591 Weiwei Guo, Xiaowei Liu, Sida Wang, Huiji Gao, Ananth Sankar, Zimeng Yang, Qi Guo, Liang
 592 Zhang, Bo Long, Bee-Chung Chen, and Deepak Agarwal. Detext: A deep text ranking frame-
 593 work with BERT. In *Proceedings of the 29th ACM International Conference on Information
 594 and Knowledge Management (CIKM '20)*, Virtual Event, Ireland, 2020. ACM. URL <https://arxiv.org/abs/2008.02460>.

595

596 Ankit Gupta, Albert Gu, and Jonathan Berant. Diagonal state spaces are as effective as structured
 597 state spaces. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2022.

594 Mohammad Hammoud and Devang Acharya. Don't pay attention. *arXiv preprint arXiv:2506.11305*,
 595 2025.

596

597 Pengcheng He, Jianfeng Gao, and Weizhu Chen. Debertav2: Further improving deberta using
 598 knowledge distillation. *arXiv preprint arXiv:2111.09543*, 2021a.

599 Pengcheng He, Xiaodong Liu, Jianfeng Gao, and Weizhu Chen. Deberta: Decoding-enhanced bert
 600 with disentangled attention. In *International Conference on Learning Representations (ICLR)*,
 601 2021b.

602 Pengcheng He, Jianfeng Gao, and Weizhu Chen. Debertav3: Improving deberta using electra-style
 603 pre-training with gradient-disentangled embedding sharing. *arXiv preprint arXiv:2111.09543*,
 604 2023.

605 Eduard Hovy, Mitch Marcus, Martha Palmer, Lance Ramshaw, and Ralph Weischedel. Ontonotes:
 606 the 90% solution. In *Proceedings of the human language technology conference of the NAACL,*
 607 *Companion Volume: Short Papers*, pp. 57–60, 2006.

608 Hugging Face. Fineweb dataset. <https://huggingface.co/datasets/HuggingFaceFW/fineweb>, 2023. Accessed: 2025-04-19.

609

610

611 Xiaoqi Jiao, Yichun Yin, Lifeng Shang, Xin Jiang, Xiao Chen, Linlin Li, Fang Wang, and Qun Liu.
 612 Tinybert: Distilling bert for natural language understanding. In *Findings of the Association for*
 613 *Computational Linguistics: EMNLP 2020*, pp. 4163–4174, 2020.

614

615 Andrej Karpathy. The most dramatic optimization to nanogpt so far (~25% speedup) is to sim-
 616 ply increase vocab size from 50257 to 50304 (nearest multiple of 64). <https://x.com/karpathy/status/1621578354024677377>, February 2023. X (formerly Twitter) post.

616

617

618 Vladimir Karpukhin, Barlas Oguz, Sewon Min, Patrick Lewis, Ledell Wu, Sergey Edunov, Danqi
 619 Chen, and Wen-tau Yih. Dense passage retrieval for open-domain question answering. In
 620 *Proceedings of the 2020 Conference on Empirical Methods in Natural Language Processing*
 621 (*EMNLP*), pp. 6769–6781, 2020.

622

623 Angelos Katharopoulos, Apoorv Vyas, Nikolaos Pappas, and François Fleuret. Transformers are
 624 rnns: Fast autoregressive transformers with linear attention. In *Proceedings of the 37th Interna-*
 625 *tional Conference on Machine Learning (ICML)*, 2020.

626

627 Omar Khattab and Matei Zaharia. Colbert: Efficient and effective passage search via contextualized
 628 late interaction over bert. In *Proceedings of the 43rd International ACM SIGIR Conference on*
 629 *Research and Development in Information Retrieval (SIGIR)*, pp. 39–48, 2020.

630

631 Tom Kwiatkowski, Jennimaria Palomaki, Olivia Redfield, Michael Collins, Ankur Parikh, Chris
 632 Alberti, Danielle Epstein, Illia Polosukhin, Jacob Devlin, Kenton Lee, et al. Natural questions: a
 633 benchmark for question answering research. *Transactions of the Association for Computational*
 634 *Linguistics*, 7:453–466, 2019.

635

636 Guokun Lai, Qizhe Xie, Hanxiao Liu, Yiming Yang, and Eduard Hovy. Race: Large-scale reading
 637 comprehension dataset from examinations. In *Proceedings of the 2017 Conference on Empirical*
 638 *Methods in Natural Language Processing (EMNLP)*, pp. 785–794. Association for Computational
 639 Linguistics, 2017.

640

641 Zhenzhong Lan, Mingda Chen, Sebastian Goodman, Kevin Gimpel, Piyush Sharma, and Radu Sori-
 642 cut. Albert: A lite bert for self-supervised learning of language representations. In *International*
 643 *Conference on Learning Representations (ICLR)*, 2020.

644

645 Shen Li, Yanli Zhao, Rohan Varma, Omkar Salpekar, Pieter Noordhuis, Teng Li, Adam Paszke,
 646 Jeff Smith, Brian Vaughan, Pritam Damania, and Soumith Chintala. Pytorch distributed: Experi-
 647 ences on accelerating data parallel training. *Proceedings of the VLDB Endowment (PVLDB)*, 13
 648 (12):3005–3018, 2020. doi: 10.14778/3407790.3407832. URL <https://www.vldb.org/pvldb/vol13/p3005-li.pdf>.

649

650 Hanxiao Liu, Zihang Dai, David So, and Quoc V Le. Pay attention to mlps. *Advances in neural*
 651 *information processing systems*, 34:9204–9215, 2021.

648 Yinhan Liu, Myle Ott, Naman Goyal, Jingfei Du, Mandar Joshi, Danqi Chen, Omer Levy, Mike
 649 Lewis, Luke Zettlemoyer, and Veselin Stoyanov. Roberta: A robustly optimized bert pretraining
 650 approach. *arXiv preprint arXiv:1907.11692*, 2019.

651

652 Stephen Mayhew, Terra Blevins, Shuheng Liu, Marek Šuppa, Hila Gonen, Joseph Marvin Imperial,
 653 Börje F Karlsson, Peiqin Lin, Nikola Ljubešić, Lester James Miranda, et al. Universal ner: A gold-
 654 standard multilingual named entity recognition benchmark. *arXiv preprint arXiv:2311.09122*,
 655 2023.

656 Niklas Muennighoff, Nouamane Tazi, Loïc Magne, Nils Reimers, Douwe Kiela, and Timo Schick.
 657 Mteb: Massive text embedding benchmark. *arXiv preprint arXiv:2307.16645*, 2023.

658

659 Multi-Linguality Multi-Functionality Multi-Granularity. M3-embedding: Multi-linguality, multi-
 660 functionality, multi-granularity text embeddings through self-knowledge distillation, 2024.

661 Tsendsuren Munkhdalai, Lingkai Kong, Tianxiao Zhang, and Jiawei Han. Infini-attention: Long-
 662 context attention in infinite sequence. *arXiv preprint arXiv:2404.07143*, 2024.

663

664 Zach Nussbaum, John X Morris, Brandon Duderstadt, and Andriy Mulyar. Nomic embed: Training
 665 a reproducible long context text embedder. *arXiv preprint arXiv:2402.01613*, 2024.

666

667 OpenAI. How to count tokens with tiktoken. https://cookbook.openai.com/examples/how_to_count_tokens_with_tiktoken, 2022. Accessed 2025-09-21.

668

669 OpenAI. tiktoken: A fast BPE tokenizer for use with openai models. <https://github.com/openai/tiktoken>, 2025. Accessed 2025-09-21.

670

671 Antonio Orvieto and Robert Gower. In search of adam’s secret sauce. *arXiv preprint arXiv:2505.21829*, May 2025. doi: 10.48550/arXiv.2505.21829. URL <https://arxiv.org/abs/2505.21829>.

672

673 Adam Paszke, Sam Gross, Francisco Massa, Adam Lerer, James Bradbury, Gregory Chanan, Trevor
 674 Killeen, Zeming Lin, Natalia Gimelshein, Luca Antiga, Alban Desmaison, Andreas Köpf, Ed-
 675 ward Yang, Zach DeVito, Martin Raison, Alykhan Tejani, Sasank Chilamkurthy, Benoit Steiner,
 676 Lu Fang, Junjie Bai, and Soumith Chintala. Pytorch: An imperative style, high-performance deep
 677 learning library. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019. URL
 678 <https://arxiv.org/abs/1912.01703>.

679

680 Bo Peng, Ruichong Zhang, Daniel Goldstein, Eric Alcaide, Xingjian Du, Haowen Hou, Jiaju Lin,
 681 Jiaxing Liu, Janna Lu, William Merrill, et al. Rwkv-7” goose” with expressive dynamic state
 682 evolution. *arXiv preprint arXiv:2503.14456*, 2025.

683

684 Jacob Portes, Alex Trott, Sam Havens, Daniel King, Abhinav Venigalla, Moin Nadeem, Nikhil
 685 Sardana, Daya Khudia, and Jonathan Frankle. Mosaicbert: A bidirectional encoder optimized for
 686 fast pretraining. *arXiv preprint arXiv:2312.17482*, 2023. doi: 10.48550/arXiv.2312.17482.

687

688 Ofir Press, Noah A. Smith, and Mike Lewis. Train short, test long: Attention with linear biases
 689 (alibi) enables input length extrapolation. *arXiv preprint arXiv:2108.12409*, 2022.

690

691 Alec Radford, Karthik Narasimhan, Tim Salimans, and Ilya Sutskever. Improving lan-
 692 guage understanding by generative pre-training. OpenAI technical report, 2018.
 693 https://cdn.openai.com/research-covers/language-unsupervised/language_understanding_paper.pdf.

694

695 Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sandhini Agarwal,
 696 Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, Gretchen Krueger, and Ilya Sutskever.
 697 Learning transferable visual models from natural language supervision. In *Proceedings of the 38th
 698 International Conference on Machine Learning (ICML)*, 2021.

699

700 Prabhakar Raghavan. How ai is powering a more helpful google, October 2020. URL <https://blog.google/products/search/search-on/>. The Keyword (Google). “BERT is
 701 now used in almost every query in English.” Accessed 2025-09-24.

702 Pranav Rajpurkar, Jian Zhang, Konstantin Lopyrev, and Percy Liang. Squad: 100,000+ questions for
 703 machine comprehension of text. In *Proceedings of the 2016 Conference on Empirical Methods in*
 704 *Natural Language Processing (EMNLP)*, 2016.

705

706 Pranav Rajpurkar, Robin Jia, and Percy Liang. Know what you don't know: Unanswerable questions
 707 for squad. In *Proceedings of the 56th Annual Meeting of the Association for Computational*
 708 *Linguistics (ACL)*, pp. 784–789, 2018.

709 Erik F Sang and Fien De Meulder. Introduction to the conll-2003 shared task: Language-
 710 independent named entity recognition. *arXiv preprint cs/0306050*, 2003.

711

712 Victor Sanh, Lysandre Debut, Julien Chaumond, and Thomas Wolf. Distilbert, a distilled version of
 713 bert: smaller, faster, cheaper and lighter. *arXiv preprint arXiv:1910.01108*, 2019.

714

715 Keshav Santhanam, Omar Khattab, Christopher Potts, and Matei Zaharia. Colbertv2: Effective and
 716 efficient retrieval via lightweight late interaction. *arXiv preprint arXiv:2112.01488*, 2022.

717

718 Noam Shazeer. Glu variants improve transformer. *arXiv preprint arXiv:2002.05202*, 2020.

719

720 Richard Socher, Alex Perelygin, Jean Wu, Jason Chuang, Christopher D Manning, Andrew Y Ng,
 721 and Christopher Potts. Recursive deep models for semantic compositionality over a sentiment
 722 treebank. In *Proceedings of the 2013 conference on empirical methods in natural language pro-*
 723 *cessing*, pp. 1631–1642, 2013.

724

725 Jianlin Su, Yu Lu, Shengfeng Pan, Ahmed Murtadha, Bo Wen, and Yunfeng Liu. Roformer: En-
 726 hanced transformer with rotary position embedding. *arXiv preprint arXiv:2104.09864*, 2021.

727

728 Yuxuan Su, Xinglin Wang, Yihan Wang, Dong Yu, Yiming Chen, Jing Lin, Tianyu Liu, Zheng
 729 Liu, Yu Chen, et al. Gte: General text embeddings with weak supervision. *arXiv preprint*
 730 *arXiv:2308.03281*, 2023.

731

732 Zhiqing Sun, Hongkun Yu, Xiaodan Song, Renjie Liu, Yiming Yang, and Denny Zhou. Mobilebert:
 733 a compact task-agnostic bert for resource-limited devices. In *Proceedings of the 58th Annual*
 734 *Meeting of the Association for Computational Linguistics (ACL)*, pp. 2158–2170, 2020.

735

736 Yi Tay, Mostafa Dehghani, Dara Bahri, and Donald Metzler. Efficient transformers: A survey. *ACM*
 737 *Computing Surveys*, 55(6):1–28, 2022.

738

739 Nandan Thakur, Nils Reimers, Johannes Daxenberger, James Cohen, and Iryna Gurevych. Beir:
 740 A heterogeneous benchmark for zero-shot evaluation of information retrieval models. In *Pro-*
 741 *ceedings of the 43rd International ACM SIGIR Conference on Research and Development in*
 742 *Information Retrieval (SIGIR)*, pp. 2325–2331, 2021.

743

744 Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N. Gomez,
 745 Łukasz Kaiser, and Illia Polosukhin. Attention is all you need. In *Advances in Neural Infor-*
 746 *mation Processing Systems (NeurIPS)*, 2017.

747

748 Alex Wang, Yada Pruksachatkun, Nikita Nangia, Amanpreet Singh, Julian Michael, Felix Hill, Omer
 749 Levy, and Samuel R. Bowman. Superglue: A stickier benchmark for general-purpose language
 750 understanding systems. In *Advances in Neural Information Processing Systems (NeurIPS)*, 2019a.

751

752 Alex Wang, Amanpreet Singh, Julian Michael, Felix Hill, Omer Levy, and Samuel R. Bow-
 753 man. Glue: A multi-task benchmark and analysis platform for natural language understand-
 754 ing. In *Proceedings of the International Conference on Learning Representations (ICLR)*, 2019b.
 755 <https://gluebenchmark.com/>.

756

757 Sheng Wang, Ming Zhao, Nguyen Bach, Tao Xu, Zhe Huang, Fei Wu, Tong Wang, Shuo Si,
 758 Zhongqiang Li, et al. Text embeddings by weakly-supervised contrastive pre-training. *arXiv*
 759 *preprint arXiv:2212.03533*, 2022.

760

761 Zhiguo Wang, Wael Hamza, and Radu Florian. Bilateral multi-perspective matching for natural
 762 language sentences. *arXiv preprint arXiv:1702.03814*, 2017.

756 Benjamin Warner, Antoine Chaffin, Benjamin Clavié, Orion Weller, Oskar Hallström, Said
 757 Taghadouini, Alexis Gallagher, Raja Biswas, Faisal Ladhak, Tom Aarsen, Griffin Thomas
 758 Adams, Jeremy Howard, and Iacopo Poli. Smarter, better, faster, longer: A modern bidirec-
 759 tional encoder for fast, memory efficient, and long context finetuning and inference. In *Pro-
 760 ceedings of the 63rd Annual Meeting of the Association for Computational Linguistics (Vol-
 761 ume 1: Long Papers)*, pp. 2526–2547, Vienna, Austria, July 2025. Association for Compu-
 762 tational Linguistics. ISBN 979-8-89176-251-0. doi: 10.18653/v1/2025.acl-long.127. URL
 763 <https://aclanthology.org/2025.acl-long.127/>.

764 Adina Williams, Nikita Nangia, and Samuel R. Bowman. A broad-coverage challenge corpus for
 765 sentence understanding through inference. In *Proceedings of the 2018 Conference of the North
 766 American Chapter of the Association for Computational Linguistics (NAACL-HLT)*, 2018.

767 Mitchell Wortsman, Peter J Liu, Lechao Xiao, Katie E Everett, Alexander A Alemi, Ben Adlam,
 768 John D Co-Reyes, Izzeddin Gur, Abhishek Kumar, Roman Novak, Jeffrey Pennington, Jascha
 769 Sohl-Dickstein, Kelvin Xu, Jaehoon Lee, Justin Gilmer, and Simon Kornblith. Small-scale prox-
 770 ies for large-scale transformer training instabilities. In *The Twelfth International Conference on
 771 Learning Representations (ICLR 2024)*, 2024. URL <https://openreview.net/forum?id=d8w0pmvXbZ>. Oral.

772

773 Lee Xiong, Chenyan Xiong, Ye Li, Kwok-Fung Tang, Jina Mao, Paul N. Bennett, Jiadong Chen,
 774 and Arnold Overwijk. Approximate nearest neighbor negative contrastive learning for dense text
 775 retrieval. In *Proceedings of the International Conference on Learning Representations (ICLR)*,
 776 2021.

777 Biao Zhang and Rico Sennrich. Root mean square layer normalization. In *Advances in Neural
 778 Information Processing Systems (NeurIPS)*, 2019.

779

780 Jeffrey Zhu. Bing delivers its largest improvement in search experience using azure
 781 gpus, November 2019. URL <https://azure.microsoft.com/en-us/blog/bing-delivers-its-largest-improvement-in-search-experience-using-azure-gpus/>.
 782 Reports serving >1M BERT inferences/sec worldwide. Accessed 2025-09-24.

783

784

785

786

787

788

789

790

791

792

793

794

795

796

797

798

799

800

801

802

803

804

805

806

807

808

809

A MONOTONICITY UNDER DECOUPLING

In Section 4.2, we claimed that decoupling static and dynamic layers maintains *monotonicity with respect to relevance* within each dynamic layer, and that this guarantee is isolated from what a subsequent static layer does. We now formalize and prove this claim.

Setup 1 (dynamic layer). As defined in Section 4.2, given $\mathbf{Z}_{\text{tr}} \in \mathbb{R}^{C \times d'}$ and fixing a target embedding (row) i , Avey-B’s dynamic layer computes:

$$\mathbf{S} = \mathcal{N}(\mathbf{Z}_{\text{tr}}) \mathcal{N}(\mathbf{Z}_{\text{tr}})^\top, \quad (9a)$$

$$\tilde{\mathbf{S}}_{i,j} = \frac{\mathbf{S}_{i,j}}{\sum_{j'=1}^C \mathbf{S}_{i,j'} + \varepsilon}, \quad (9b)$$

$$\mathbf{c}_{\text{dyn},i} = \sigma \left(\sum_{j=1}^C \tilde{\mathbf{S}}_{i,j} \mathbf{Z}_{\text{tr},j} + \mathbf{b}_i^{(d)} \right), \quad (9c)$$

where $\mathcal{N}(\cdot)$ is per-row ℓ_2 normalization (so \mathbf{S} contains cosine similarities), $\varepsilon > 0$ is a stabilizer, σ is a pointwise monotone activation, and $\mathbf{b}_i^{(d)}$ is an optional bias.

Assumptions. We assume the following:

(A1) Nonnegative similarities. The enricher uses a nonnegative pointwise activation, namely, ReLU² (Hammoud & Acharya, 2025), hence, rows of $\mathcal{N}(\mathbf{Z}_{\text{tr}})$ are nonnegative, implying $\mathbf{S}_{i,j} \geq 0$.

(A2) Positive normalization. For each row i , $\sum_{j=1}^C \mathbf{S}_{i,j} + \varepsilon > 0$ with $\varepsilon > 0$.

810 (A3) **Monotone activation.** σ is monotone nondecreasing (e.g., Avey-B uses ReLU in Equation 9c).
 811

812 **Proposition A.1** (dynamic layer monotonicity). For a fixed target row i and any two embeddings
 813 $j_1, j_2 \in \{1, \dots, C\}$:

814 (i) *Order preservation.* If $\mathbf{S}_{i,j_1} \geq \mathbf{S}_{i,j_2}$ then $\tilde{\mathbf{S}}_{i,j_1} \geq \tilde{\mathbf{S}}_{i,j_2}$.

815 (ii) *Self-monotonicity.* Increasing $\mathbf{S}_{i,j}$ (while holding $\{\mathbf{S}_{i,k}\}_{k \neq j}$ fixed) weakly increases $\tilde{\mathbf{S}}_{i,j}$
 816 and does not increase $\tilde{\mathbf{S}}_{i,k}$ for $k \neq j$.

817 Consequently, a more relevant token (higher similarity) receives at least as large (and typically
 818 larger) weight than a less relevant token, and increasing its relevance cannot reduce or flip the sign
 819 of its contribution in the update within the dynamic layer (i.e., the update is monotone with respect
 820 to relevance).

821 *Proof.* Let $d_i = \sum_{j=1}^C \mathbf{S}_{i,j} + \varepsilon$; by (A2), $d_i > 0$. Since $\tilde{\mathbf{S}}_{i,j} = \mathbf{S}_{i,j}/d_i$ for fixed i , dividing by a
 822 positive constant preserves order. In addition, treating d_i as a function of $\mathbf{S}_{i,\cdot}$ and using (A1):
 823

$$824 \frac{\partial \tilde{\mathbf{S}}_{i,j}}{\partial \mathbf{S}_{i,j}} = \frac{d_i - \mathbf{S}_{i,j}}{d_i^2} \geq 0, \quad \text{since } d_i \geq \mathbf{S}_{i,j} \text{ by (A1) and } \varepsilon > 0, \quad (10a)$$

$$825 \frac{\partial \tilde{\mathbf{S}}_{i,k}}{\partial \mathbf{S}_{i,j}} = -\frac{\mathbf{S}_{i,k}}{d_i^2} \leq 0, \quad k \neq j, \quad \text{by (A1).} \quad (10b)$$

826 Thus, increasing a token’s similarity weakly increases (or cannot reduce) its own normalized weight
 827 (by Equation 10a) and weakly decreases (or does not increase) others’ (by Equation 10b). By (A1),
 828 $\tilde{\mathbf{S}}_{i,j} \geq 0$, so each token’s influence enters the pre-activation with a nonnegative coefficient; by
 829 (A3), σ cannot invert these contributions. Hence, the dynamic update is monotone with respect to
 830 relevance. \blacksquare

831 **Setup 2 (static layer).** Consider a dynamic layer at depth ℓ followed by a static layer:

$$832 \mathbf{h}^{(\ell+1)} = \sigma(\tilde{\mathbf{S}}^{(\ell)} \mathbf{Z}_{\text{tr}}^{(\ell)} + \mathbf{b}^{(d)}), \quad (11)$$

$$833 \mathbf{y}^{(\ell+2)} = \sigma(\mathbf{V} \mathbf{h}^{(\ell+1)} + \mathbf{b}^{(s)}), \quad (12)$$

834 where $(\mathbf{V}, \mathbf{b}^{(s)})$ are learned parameters that do not depend on $\mathbf{S}^{(\ell)}$ or $\tilde{\mathbf{S}}^{(\ell)}$.

835 **Proposition A.2** (static layer non-violation). A static layer (as in equation 12) cannot violate the
 836 monotonicity guarantee of Proposition A.1 established for a preceding dynamic layer (as in Equa-
 837 tion 11) at depth ℓ .
 838

839 *Proof.* The monotonicity statements in Proposition A.1 concern only the relationship between
 840 the relevance scores $\mathbf{S}^{(\ell)}$ and the normalized scores $\tilde{\mathbf{S}}^{(\ell)}$ used *inside* the dynamic update equa-
 841 tion 11. The static map $\mathbf{h}^{(\ell+1)} \mapsto \mathbf{y}^{(\ell+2)}$ depends on $\mathbf{h}^{(\ell+1)}$ and the similarity-agnostic parameters
 842 $(\mathbf{V}, \mathbf{b}^{(s)})$; it neither accesses nor alters $\mathbf{S}^{(\ell)}$ or $\tilde{\mathbf{S}}^{(\ell)}$. Therefore, composing the dynamic update
 843 with a static layer cannot change the inequalities and partial orders that define monotonicity for the
 844 dynamic layer’s scores. \blacksquare

845 **Remark.** A static layer reshapes representations that subsequent dynamic layers will use to com-
 846 pute new similarities, but it does not retroactively modify the scores already assigned by a preceding
 847 dynamic layer. Thus, monotonicity holds at each dynamic layer, and decoupling preserves this
 848 guarantee throughout the stack.

864
865
866 Table 3: A comparison of all the evaluated encoders across different dimensions.
867
868
869
870
871
872
873
874
875

Dimension	BERT		RoBERTa		ModernBERT		NeoBERT medium	Avey-B	
	base	large	base	large	base	large		base	large
Parameters	120M	350M	125M	355M	149M	395M	250M	165M	391M
Data Sources	BooksCorpus Wikipedia		BooksCorpus OpenWebText Stories / CC-News		Undisclosed		RefinedWeb		FineWeb
Pre-training Context Width	512		512		1,024 → 8,192		1,024 → 4,096		2,048
Inference Sequence Length	512		512		8,192		4,096		∞
Masking Rate	15%		15%		30%		20%		20%
Masking Scheme	80/10/10		80/10/10		—		100		100
Tokens Seen	131B		131B		~2T		2.1T		180B

876
877
878 **B PRETRAINING METHODOLOGY**
879

880 In this section, we detail the pretraining setup. For Avey-B, we adopt the same tokenizer
881 as Avey (Hammoud & Acharya, 2025), namely, a BPE tokenizer derived from OpenAI’s
882 p50k_base (OpenAI, 2022; 2025), with the vocabulary size set to 50,304 to align with multiple-
883 of-64 boundaries and improve hardware efficiency (Karpathy, 2023). We retain BERT-style special
884 tokens for backward compatibility with downstream applications, while using only the [MASK]
885 token during pre-training.

886 We pretrain two Avey-B sizes, *base* (165M) and *large* (391M), for 180B tokens drawn from
887 the FineWeb 300BT split (Hugging Face, 2023), using PyTorch DDP across 16 NVIDIA H200
888 GPUs (Paszke et al., 2019; Li et al., 2020). The global batch size is set to 512K tokens for both
889 models and we utilize the AdamW optimizer with $\beta_1 = \beta_2 = 0.95$ (Orvieto & Gower, 2025),
890 $\epsilon = 10^{-18}$ (Wortsman et al., 2024), weight decay of 0.01, and gradient clipping at 1.0. For the
891 learning-rate schedule, we employ a 10% linear warmup to 5×10^{-4} (base) or 2.5×10^{-4} (large),
892 followed by cosine decay to zero over the remaining 90% of steps.

893 For ablations and design-choice studies, we use the Avey-B *base* model, pretrained with a constant
894 learning rate of 10^{-3} for 10B tokens. During pretraining, sequences are *packed* so that each training
895 example meets the target sequence length, following the original Avey setup. We train with a masked
896 language modeling (MLM) objective, randomly masking 20% of tokens per example after exploring
897 several masking rates (see Appendix G).

898 Finally, for the Transformer-based encoders, we use publicly available pretrained checkpoints
899 from the Hugging Face Hub (google-bert-base; google-bert-large; facebook-roberta-base; facebook-
900 roberta-large; answerdotai-base; answerdotai-large; chandar-lab). Table 3 summarizes the evaluated
901 models along key dimensions, including parameter count, context window, and pretraining tokens,
902 among others.

903
904 **C SHOULD THE RANKER OPERATE BIDIRECTIONALLY?**
905

906 In the original, unidirectional Avey architecture, the ranker attends only to *preceding* splits (left
907 context) to preserve the causal constraint of autoregressive modeling. With bidirectional Avey-B,
908 we ask whether the ranker should, like the neural processor, operate bidirectionally, retrieving from
909 both left and right contexts of the current split.

910 To this end, we conduct all our experiments using the Avey-B *base* variant (165M parameters) (see
911 Table 3). We pretrain it on 10B FineWeb (Hugging Face, 2023) tokens with a constant learning
912 rate of 1×10^{-3} and a 20% masking rate. As suggested in Section 5.1, we finetune SC and TC
913 for one epoch, QA for four epochs, and IR for 1,000 optimization steps. For each task, we run 5
914 independent seeds with a learning rate of 5×10^{-4} , using a 10% warmup followed by linear decay
915 to zero over the remaining 90% of steps. We report the *best-of-5* score for each configuration as
916 an upper bound. Metrics are accuracy for SC, F1 for TC and QA, and NDCG@10 for IR. Table 4
917 shows all the results.

918 Table 4: Effectiveness results comparing unidirectional vs. bi-directional rankers.
919

920 Ranker Type	921 SC				921 TC				921 QA				921 IR			
	MNLI	QQP	SST-2	Avg.	CONLL	Onto.	UNER	Avg.	ReCoRD	SQuAD	SQuADv2	Avg.	MLDR	MS	MARCO	NQ
922 Unidirectional ranker	81.39	88.92	91.86	87.39	92.96	93.21	93.97	93.38	30.22	62.26	60.72	51.07	60.27	87.48	76.71	74.82
923 Bi-directional ranker	80.54	88.45	91.74	86.91	92.62	92.91	94.35	93.29	10.64	45.34	53.56	36.51	60.05	90.02	48.54	66.20

924
925 In particular, Table 4 contrasts a *unidirectional* ranker with a *bidirectional* ranker across SC, TC, QA,
926 and IR task categories. As illustrated, the bidirectional ranker consistently underperforms
927 the unidirectional ranker. While the gap is modest on SC (87.39 → 86.91; $\Delta = -0.48$) and TC
928 (93.38 → 93.29; $\Delta = -0.09$), it is substantial on QA (51.07 → 36.51; $\Delta = -14.56$) and IR
929 (74.82 → 66.20; $\Delta = -8.62$). Notably, QA performance degrades sharply, with F1 of ReCoRD
930 dropping from 30.22 to 10.64, suggesting that right-context retrieval at the split level may severely
931 harm evidence selection for reasoning.

932 We identify two likely causes for this behavior. First, a unidirectional ranker enforces causal ordering
933 and encourages the model to accumulate evidence along the discourse flow. Natural language
934 often exhibits forward dependencies, whereby content in a later split is best interpreted in light
935 of earlier splits. As such, allowing the current split to pair with future splits can dilute or over-
936 ride strong signals from its relevant preceding splits. Second, Avey-B already provides look-ahead,
937 *token-level* contextualization within each split (the contextualizer operates without a causal mask on
938 each split) so every position in it except the last has access to rightward tokens. Therefore, additional
939 look-ahead, *split-level* contextualization seems often redundant and at times even disruptive.

940 In summary, these findings suggest that while look-ahead, token-level contextualization within a
941 split benefits Avey-B, look-ahead, split-level contextualization, driven by the ranker attending to
942 both left and right split contexts of the current split, is not advantageous and potentially counterpro-
943 ductive.

944 D HOW TO ARRANGE STATIC AND DYNAMIC LAYERS?

945 Decoupling static and dynamic parameterizations into distinct layer types in Avey-B introduces a
946 key architectural degree of freedom, that is, *how to arrange static (S) and dynamic (D) layers across*
947 *depth*. We therefore evaluate the following families of patterns and report their effectiveness:

- 950 **1. Interleaved:** Alternate $S \leftrightarrow D$. With an even number of layers, we test both start points,
951 $S \rightarrow D \rightarrow \dots$ and $D \rightarrow S \rightarrow \dots$
- 953 **2. Single dynamic:** Exactly one D and the remainder S , placing D either at the *head* (to prime
954 downstream static transformations) or at the *tail* (to refine final representations).
- 955 **3. Two-stage stack:** First half one type and second half the other, considering both orders ($S^{L/2} \rightarrow$
956 $D^{L/2}$ and $D^{L/2} \rightarrow S^{L/2}$).
- 957 **4. Uniform stack:** Either all-static or all-dynamic stack, as boundary conditions.

958 We utilize the experimental setup described in Appendix C. Table 5 reports all the results across
959 SC, TC, QA, and IR task categories. Two consistent trends emerge. First, the *interleaved* ar-
960 rangement, $S \rightarrow D \rightarrow \dots$, attains the strongest average performance on SC, TC, and QA, while
961 remaining competitive on IR. This suggests that a static front layer is potentially providing a sta-
962 ble representational “scaffold” before any input-dependent mixing, reducing variance introduced by
963 raw similarity scores and improving downstream contextualization. Second, the interleaved pattern,
964 $D \rightarrow S \rightarrow \dots$, underperforms the $S \rightarrow D \rightarrow \dots$ variant (most notably on QA and IR) likely because
965 early, similarity-driven updates are fragile without a learned (static) basis to shape features prior to
966 dynamic contextualization.

967 The *uniform stack, all-static* configuration performs worse than interleaved arrangements but only
968 modestly so (it even slightly outperforms them on IR), indicating that static linear projections alone
969 already enable strong contextualization, even without any input-dependent adaptation. Conversely,
970 the *uniform stack, all-dynamic* pattern performs worse across all benchmark categories (particularly
971 QA). The *single-dynamic* and *two-stage stack* arrangements fall between these extremes, though

972
973

Table 5: Effectiveness results across different static (S) and dynamic (D) layering patterns.

Pattern	SC				TC				QA				IR			
	MNLI	QQP	SST-2	Avg.	CONLL	Onto.	UNER	Avg.	ReCoRD	SQuAD	SQuADv2	Avg.	MLDR	MS	MARCO	NQ
Interleaved, $S \rightarrow D \rightarrow \dots$	81.39	88.92	91.86	87.39	92.96	93.21	93.97	93.38	30.22	62.26	60.72	51.07	60.27	87.48	76.71	74.82
Interleaved, $D \rightarrow S \rightarrow \dots$	77.89	87.51	90.37	85.26	91.73	92.48	93.24	92.48	21.31	56.93	55.77	44.67	52.67	88.47	68.63	69.92
Single dynamic as a head	73.69	86.55	91.06	83.77	92.42	92.75	93.31	92.83	22.42	56.94	54.87	44.74	60.22	89.70	73.32	74.41
Single dynamic as a tail	73.19	87.90	91.63	84.24	92.58	93.10	93.80	93.16	24.64	55.29	54.06	44.66	60.55	87.72	75.31	74.53
Two-stage stack, $S^{L/2} \rightarrow D^{L/2}$	75.70	86.56	90.25	84.17	92.51	92.65	94.24	93.13	15.29	52.18	55.09	40.85	54.97	85.92	67.54	69.48
Two-stage stack, $D^{L/2} \rightarrow S^{L/2}$	74.37	87.15	91.17	84.23	92.70	93.03	93.72	93.15	29.30	54.36	51.20	44.95	59.28	89.56	76.74	75.19
Uniform stack, all-static	77.58	87.97	91.51	85.69	92.66	93.10	94.04	93.27	23.85	56.43	54.73	45.00	62.54	86.38	75.92	74.95
Uniform stack, all-dynamic	68.04	83.26	87.27	79.52	90.31	90.86	91.17	90.78	19.49	47.51	50.44	39.15	57.70	89.33	69.50	72.18

981

982

Table 6: Effectiveness results across different normalization schemes.

Normalization Scheme	SC				TC				QA				IR			
	MNLI	QQP	SST-2	Avg.	CONLL	Onto.	UNER	Avg.	ReCoRD	SQuAD	SQuADv2	Avg.	MLDR	MS	MARCO	NQ
Divide-by-sum norm	81.39	88.92	91.86	87.39	92.96	93.21	93.97	93.38	30.22	62.26	60.72	51.07	60.27	87.48	76.71	74.82
RMS norm	64.70	87.21	88.76	80.22	90.82	92.12	91.89	91.61	21.37	56.09	56.33	44.60	50.27	89.33	66.68	68.76
Softmax	79.31	88.16	91.06	86.18	92.39	92.96	93.45	92.93	27.70	59.29	58.55	48.51	61.83	89.75	74.43	75.34
Scaled softmax	76.70	87.24	91.86	85.27	92.63	93.02	94.07	93.24	24.14	58.79	56.23	46.39	62.24	87.63	74.39	74.75

989

990

they typically trail the interleaved static-first design (except for two-stage stack, $D^{L/2} \rightarrow S^{L/2}$ on IR). Overall, these findings highlight that while dynamic parameterization contributes meaningfully to performance, it is most effective when interleaved with static layers that supply a stable basis and representational depth.

994

995

E HOW TO NORMALIZE?

997

In Avey-B, dynamic layers contextualize tokens by constructing a cosine similarity matrix from pairwise cosine scores of the input. Since the similarity scores are used to perform a weighted sum of input embeddings at every position and the sum of the raw similarity magnitudes can vary significantly, we tested several normalization strategies to stabilize training and improve generalization, including *divide-by-sum norm* (i.e., row-wise normalization by the sum of similarities), *RMS norm* (i.e., row-wise normalization by root mean square), *softmax*, and *scaled softmax* (with temperature scaling, analogous to scaled dot-product attention). In all the tests, we used the same experimental setup discussed in Appendix C.

1005

As shown in Table 6, the simple divide-by-sum norm method achieves the strongest overall performance, outperforming alternatives on SC, TC, and QA, and almost matching or surpassing them on IR. Notably, divide-by-sum norm provides a balanced distribution of contextual weights while retaining sign information, which is lost under softmax-based schemes. By contrast, softmax and scaled softmax yield weaker SC, TC, and QA scores but softmax outperforms divide-by-sum norm on IR. On average, RMS norm underperforms divide-by-sum norm across all categories.

1011

These findings indicate that unlike self-attention, which benefits from exponential normalization (as provided by softmax), Avey-B’s cosine-based dynamic layers benefit from a *conservative, structure-preserving* normalization. Exponentiation amplifies outliers, distorts relative similarity ratios, and can swamp the static path. In contrast, *divide-by-sum norm* preserves the ordering and margins of similarities, constrains each row to a convex combination (weights in $[0, 1]$ that sum to 1), and effectively bounds the operator norm, yielding stable gradients and preventing the dynamic stream from overwhelming the static contributions. Empirically, this simple choice delivers strong gains across SC, TC, QA, and IR while maintaining robust training dynamics.

1019

1020

F WHAT ARE THE BEST SEQUENCE LENGTH, SPLIT SIZE, AND TOP- k VALUES?

1023

1024

We now analyze how Avey-B’s downstream performance is affected by the ranker’s three hyperparameters, namely, the training sequence length N , the split size S , and the number of top- k splits selected for contextualization. N governs the size of the candidate pool available to the ranker, S

Table 7: Effectiveness results across different sequence length N , split size S , and top- k values.

N	S	k	SC				TC				QA				IR								
			MNLI			Avg.	CONLL			Onto.	UNER	Avg.	ReCoRD			SQuAD	SQuADv2	Avg.	MLDR	MS	MARCO	NQ	Avg.
			MNLI	QQP	SST-2	Avg.	CONLL	Onto.	UNER	Avg.	ReCoRD	SQuAD	SQuADv2	Avg.	MLDR	MS	MARCO	NQ	Avg.	MS	MARCO	NQ	Avg.
512	128	1	80.09	88.82	91.63	86.85	92.69	93.08	93.85	93.21	27.65	60.44	58.82	48.97	60.06	88.89	76.11	75.02	512	1024	1024	1024	1024
		3	80.95	88.89	91.63	87.16	92.74	93.03	93.55	93.11	28.11	60.66	58.79	49.19	60.72	90.11	76.06	75.63					
		256	1	79.68	88.75	92.32	86.92	92.27	92.94	93.69	92.97	43.70	71.25	64.06	59.67	60.71	88.61	76.29	75.20				
1024	128	1	80.79	88.80	91.63	87.07	92.87	93.03	93.81	93.24	27.42	61.14	58.85	49.14	60.44	88.95	76.05	75.15	1024	1024	1024	1024	1024
		3	80.02	88.71	91.63	86.79	92.89	93.06	94.43	93.46	28.74	61.55	60.49	50.26	61.17	90.26	75.66	75.70					
		5	80.64	88.62	91.40	86.89	92.86	93.13	93.88	93.29	28.69	61.49	59.79	49.99	61.81	86.86	75.61	74.76					
1024	256	7	80.70	89.02	91.74	87.15	92.44	93.14	93.83	93.14	27.96	63.47	59.98	50.47	62.58	86.40	66.73	71.90	1024	1024	1024	1024	1024
		1	79.53	88.80	91.28	86.54	92.10	93.02	93.57	92.90	42.82	70.16	62.54	58.51	60.88	88.87	76.34	75.36					
		3	79.95	88.65	91.40	86.67	92.10	92.84	93.48	92.81	42.51	70.55	63.16	58.74	59.14	89.43	77.05	75.21					
1024	512	1	79.41	88.35	91.51	86.42	92.65	93.18	93.74	93.19	43.96	71.92	63.91	59.93	60.26	89.20	77.82	75.76	1024	1024	1024	1024	1024
		3	80.60	88.93	91.74	87.09	92.85	93.13	94.05	93.34	28.72	61.50	59.13	49.78	60.21	90.69	75.15	75.35					
		5	80.94	88.95	91.63	87.17	92.52	93.00	93.68	93.07	28.49	60.95	59.73	49.72	60.71	89.22	75.48	75.14					
1024	128	7	80.98	88.66	91.51	87.05	92.86	93.17	94.02	93.35	27.07	60.27	59.23	48.86	59.02	88.43	75.73	74.39	1024	1024	1024	1024	1024
		9	67.24	88.42	91.63	82.43	91.77	92.69	93.00	92.49	27.76	61.81	58.45	49.34	57.08	86.89	73.55	72.51					
		11	80.64	88.84	91.51	87.00	92.54	92.97	93.82	93.11	29.82	62.44	59.52	50.59	60.77	88.27	75.93	74.99					
1024	2048	13	81.14	88.75	92.20	87.36	92.92	93.13	94.05	93.37	27.36	59.98	58.01	48.45	60.21	90.31	78.48	76.33	1024	1024	1024	1024	1024
		15	80.78	89.01	91.63	87.14	92.51	92.94	94.11	93.19	28.27	60.34	59.46	49.36	58.69	88.11	76.03	74.28					
		1	70.61	88.69	92.09	83.80	92.38	93.18	94.02	93.19	44.32	70.96	63.58	59.62	60.84	88.15	77.24	75.41					
1024	256	3	80.18	88.91	91.74	86.94	92.33	93.05	93.75	93.04	42.91	71.99	64.81	59.90	60.02	90.36	76.58	75.65	1024	1024	1024	1024	1024
		5	79.45	88.48	90.71	86.21	91.51	92.46	92.77	92.25	38.85	69.88	62.09	56.94	56.72	88.11	73.70	72.84					
		7	78.98	88.65	91.74	86.46	93.00	92.85	93.66	93.17	42.89	70.64	63.25	58.93	59.42	88.35	76.36	74.71					
1024	512	1	79.37	88.52	91.63	86.51	92.41	92.77	93.26	92.81	41.67	71.45	62.20	58.44	56.76	88.67	77.94	74.46	1024	1024	1024	1024	1024
		3	79.36	88.67	91.28	86.44	92.15	92.87	94.00	93.01	46.39	72.34	64.06	60.93	55.65	87.32	75.89	72.95					
		1024	1	75.29	88.34	91.40	85.01	92.10	92.79	93.57	92.82	44.88	70.47	62.05	59.13	55.86	90.07	73.93	73.29				

determines the size (in tokens) of each candidate split, and the effective context width seen by the contextualizer is $C = S(k+1)$. We follow the experimental setup described in Appendix C. Table 7 illustrates all the results.

To begin with, the dominant trend across tasks is that performance peaks when the effective context $C = S(k+1)$ matches or closely approximates the training sequence length N . For example, on QA with $N=2048$, the best average occurs at $S=512$, $k=3$, giving $C=512 \times (3+1)=2048=N$. For SC, TC, and IR at $N=2048$, the strongest averages are at $S=128$, $k=13$, yielding $C=128 \times (13+1)=1792$, close to N . Similar behavior holds for $N=512$ and $N=1024$ across categories, with one slight exception, that is, TC. In particular, on TC, the best setting often lands on $C \approx N/2$ (e.g., at $N=512$ the optimum is $S=128$, $k=1$, so $C=256=N/2$, but it is only +0.1 points away from the effectiveness at $N=512$, $S=128$, $k=3$, which yields $C=128 \times (3+1)=512=N$).

In summary, Avey-B’s performance generally improves with a larger training sequence length N . In our experiments, the best results occur at the largest tested $N=2048$ for SC, QA, and IR. The exception is TC, which peaks at $N=512$ but is within +0.09 points of the $N=1024$ setting. Across these optima, the coverage heuristic $C = S(k+1) \approx N$ is consistently satisfied (matching or closely approaching N). This pattern suggests that, for a bidirectional encoder, one should enlarge the candidate pool via larger N while ensuring ample contextual coverage by setting $S(k+1)$ to match or closely track N . Averaging over all task categories, the best overall configuration is $N=2048$, $S=256$, $k=3$, hence, it was adopted as Avey-B’s default configuration.

G WHAT IS THE BEST MASKING RATE?

Because Avey-B is pretrained with masked language modeling, the fraction of tokens replaced by the [MASK] token sets the task difficulty. In particular, too little masking makes reconstruction nearly trivial, whereas too much masking deprives the model of sufficient contextual signal for reliable prediction. To calibrate this trade-off, we swept masking rates from 10% to 50% for both the *base* and *large* models (see Table 3 in Appendix B), while following the same experimental setup described in Appendix C (except for the masking rate since we vary it here).

1080
1081
1082 Table 8: Effectiveness results at different masking rates for Avey-B’s *base* model.
1083
1084
1085
1086
1087

Masking %	SC				TC				QA				IR			
	MNLI	QQP	SST-2	Avg.	CONLL	Onto.	UNER	Avg.	ReCoRD	SQuAD	SQuADv2	Avg.	MLDR	MS	MARCO	NQ
10%	78.07	88.02	91.86	85.98	91.61	92.54	93.50	92.55	36.64	68.57	60.65	55.29	51.02	85.75	71.50	69.42
20%	80.18	88.91	91.74	86.94	92.33	93.05	93.75	93.04	42.91	71.99	64.81	59.90	60.02	90.36	76.58	75.65
30%	78.62	88.49	91.97	86.36	92.45	93.01	93.75	93.07	42.80	71.26	63.85	59.30	62.16	89.72	76.19	76.02
40%	77.05	88.02	91.51	85.53	92.26	92.90	93.49	92.88	39.70	69.84	62.45	57.33	59.67	90.33	74.56	74.85
50%	66.12	88.32	91.06	81.83	92.86	92.62	93.16	92.88	42.15	70.44	62.90	58.50	62.03	90.06	77.69	76.59

1088
1089 Table 9: Effectiveness results at different masking percentages for Avey-B’s *large* model.
1090
1091

Masking %	SC				TC				QA				IR			
	MNLI	QQP	SST-2	Avg.	CONLL	Onto.	UNER	Avg.	ReCoRD	SQuAD	SQuADv2	Avg.	MLDR	MS	MARCO	NQ
10%	81.01	88.54	92.09	87.21	92.39	92.91	93.32	92.87	42.46	71.53	64.63	59.54	54.58	88.73	75.25	72.85
20%	82.12	89.19	92.32	87.88	92.76	92.94	93.65	93.12	47.93	72.84	65.79	62.19	63.53	91.54	80.47	78.51
30%	81.54	89.43	91.74	87.57	92.59	92.98	93.83	93.13	48.16	73.44	66.52	62.71	61.72	89.31	81.76	77.60
40%	70.21	89.11	92.20	83.84	92.51	93.13	93.65	93.10	46.29	73.36	65.71	61.79	64.09	90.96	81.35	78.80
50%	78.19	89.02	92.20	86.47	92.83	92.95	93.83	93.20	46.99	71.28	63.89	60.72	61.63	90.49	79.37	77.16

1097
1098 As shown in Table 8, increasing the masking rate from 10% to 20% improves performance across
1099 SC, TC, QA, and IR for the base model. Overall, scores typically peak at around 20%–30% masking,
1100 yielding consistent gains on SC, TC, and QA. The exception is IR, which attains its best results at
1101 50% masking. At higher masking levels (40%–50%), performance can drop markedly (e.g., MNLI),
1102 indicating that the smaller-capacity model struggles when too little context is visible (masking be-
1103 comes overly aggressive, weakening both the input signal and the training target).

1104 The larger model is more robust to masking but still follows a similar trend to the base variant
1105 (see Table 9). Performance generally improves from 10% to 20–30% masking, which offers the
1106 best cross-task trade-off (an exception is TC, which peaks at 50%). QA and IR benefit the most,
1107 with ReCoRD, MLDR, and NQ rising by over +5, +9, and +6 points, respectively, relative to 10%
1108 masking. Although the large model tolerates 40–50% masking with modest degradation, SC remains
1109 sensitive, whereby at 40% masking, MNLI drops by ∼12 points versus 20%, then partially recovers
1110 at 50%, going up by ∼8 points versus 40%. These patterns indicate that overly aggressive masking
1111 can destabilize training even at higher capacity.

1112 Overall, both models indicate 20–30% masking as near-optimal, thus, we pretrain both at 20%.

1114 H ABLATION STUDY

1116 In this study, we conduct a series of ablation experiments on Avey-B. To this end, we fix (1) the
1117 sequence length N , split size S , and top- k retrieval depth to the best settings from Appendix F;
1118 (2) the static–dynamic interleaving pattern to the best arrangement from Appendix D; and (3) the
1119 dynamic-layer normalization to the most effective scheme from Appendix E. In addition, we follow
1120 the experimental setup described in Appendix C.

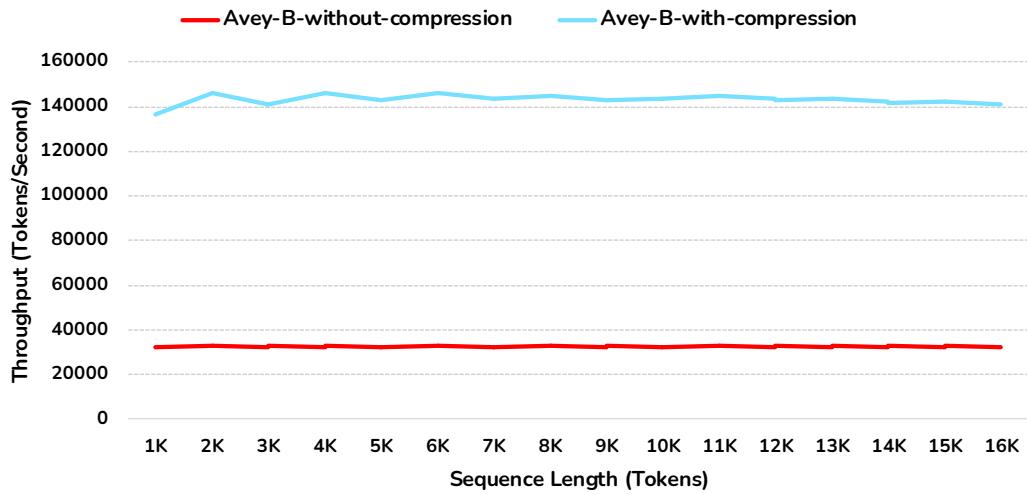
1122 Table 10 reports ablations over five key architectural components of Avey-B: (1) decoupling static
1123 and dynamic parameterizations; (2) applying row-wise normalization in the dynamic layers; (3)
1124 incorporating a neural compressor within the ranker; (4) adding a residual connection between the
1125 compressor output and the original tokens of the current split that is being compressed with its top- k
1126 retrieved splits; and (5) removing the ranker entirely.

1127 As illustrated in Table 10, coupling the static and dynamic parameterizations (i.e., see “full design”
1128 and “w/o decoupling” rows) yields consistent drops of 1.43%, 2.12%, 2.53%, and 7.40% on SC,
1129 TC, QA, and IR, respectively, confirming that separating similarity scoring from neural learning
1130 improves accuracy. Row-wise normalization proves even more critical, whereby removing it leads
1131 to significantly larger degradations of 3.55%, 0.87%, 7.65%, and 15.33% across SC, TC, QA, and
1132 IR, respectively.

1133 As discussed in Section 4.3, the neural compressor reduces the number of contextualized tokens
1134 per split from $(k+1)S$ to S . This reduction yields a substantial $4.37\times$ throughput improvement

1134 Table 10: Ablations of Avey-B, removing one component at a time while holding all others fixed.
 1135 (1) *w/o normalization*: removes row-wise normalization in the dynamic layers; (2) *w/o decoupling*:
 1136 reverts to coupled static and dynamic parameterizations; (3) *w/o compression*: omits the neural
 1137 compressor; (4) *w/o residual*: discards the residual connection between the compressor output and
 1138 the current split’s tokens; and (5) *w/o ranker*: disables the ranker entirely.

Model	SC				TC				QA				IR			
	MNLI	QQP	SST-2	Avg.	CONLL	Onto.	UNER	Avg.	ReCoRD	SQuAD	SQuAD	Avg.	MLDR	MS	NQ	Avg.
Avey-B (full design)	80.74	88.91	91.97	87.20	91.84	93.25	93.09	92.72	39.60	68.52	60.48	56.20	57.49	90.38	75.64	74.50
Avey-B w/o normalization	77.47	84.60	90.25	84.10	90.98	92.65	92.12	91.91	29.72	67.15	58.83	51.90	46.43	82.89	59.92	63.08
Avey-B w/o decoupling	79.94	88.60	89.33	85.95	89.10	92.11	91.06	90.75	36.86	67.89	59.57	54.77	55.00	82.77	69.18	68.98
Avey-B w/o compression	80.80	89.03	91.17	87.00	91.52	93.29	92.97	92.59	42.80	70.54	59.77	57.70	60.95	89.13	76.92	75.66
Avey-B w/o residual	77.53	87.48	90.37	85.12	90.80	92.44	91.17	91.47	34.71	66.08	56.02	52.27	55.22	86.74	71.72	71.22
Avey-B w/o ranker	77.36	87.32	88.76	84.48	90.20	92.39	90.74	91.11	25.85	66.27	57.92	50.01	38.28	85.35	61.93	61.85

Figure 3: The throughput of Avey-B *with* and *without* the neural compressor.

(see Fig. 3) while preserving strong task performance (see Table 10). On SC and TC, compression has negligible effect and even produces slight average gains of +0.23% and +0.14%, respectively, likely due to the removal of noisy global tokens that can arise under top-k retrieval¹. In contrast, QA and IR exhibit modest average drops of 2.68% and 1.56%, respectively. These tasks rely more heavily on fine-grained cross-split evidence and subtle retrieval cues, which compression may partially attenuate. Overall, considering the 4.37 \times efficiency improvement, the negligible (and sometimes positive) impact on SC and TC, and the modest reductions on QA and IR, the neural compressor provides a clear and favorable efficiency–effectiveness trade-off.

Besides compression, Avey-B adds a residual connection between the compressor output and the current split’s S tokens to preserve local signal. Table 10 shows that removing this residual degrades every benchmark, with an average reduction of 3.38%, underscoring its role in maintaining lexical fidelity and stabilizing context integration.

Finally, as described in (Hammoud & Acharya, 2025), the ranker is invoked *only once* per forward/backward pass, prior to the first layer of the neural processor, and retrieves the top- k relevant splits for each current split using shallow (initial) embeddings. One might expect that deeper contextualized embeddings could yield better retrieval; however, removing the ranker entirely (i.e., allowing the neural processor to fit and operate on the entire sequence) results in universal degradation across all benchmarks, with a large average drop of 7.46% (see Table 10). This corroborates that retrieval is essential for Avey-B’s effectiveness.

¹With a split containing S tokens and the absence of a hard relevance threshold (where the ranker retrieves the top- k splits for each query split regardless of their absolute relevance), it is possible for some retrieved splits to contain weakly relevant or noisy tokens.

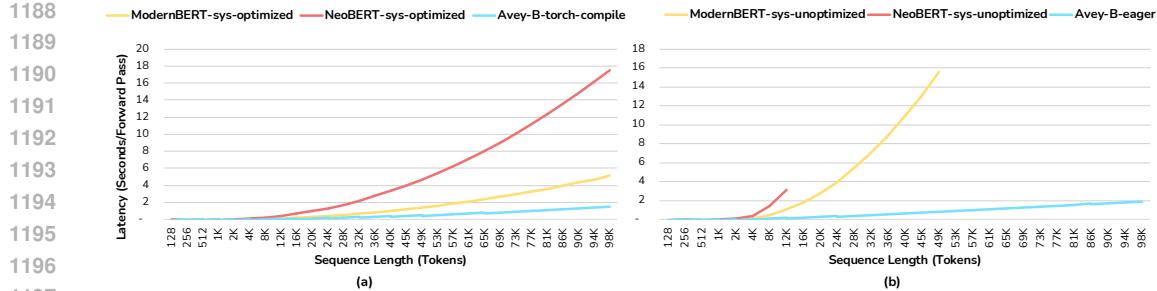


Figure 4: Latency of Avey-B, ModernBERT, and NeoBERT on NVIDIA B200 GPUs with mixed precision (BF16). We use Avey-B *base*, ModernBERT *base*, and NeoBERT *medium* (the only publicly available size). Avey-B is shown in (a) as optimized using `torch.compile` (no fused-kernel implementation is available yet) and in (b) as unoptimized (eager). For ModernBERT and NeoBERT, latency is shown for system-optimized (*with* FlashAttention) and system-unoptimized (eager) variants in (a) and (b), respectively.

Importantly, retrieval based on shallow or static embeddings is a standard practice in dense retrieval and matching systems such as DPR (Karpukhin et al., 2020), ColBERT (Khattab & Zaharia, 2020), ANCE (Xiong et al., 2021), and CLIP (Radford et al., 2021), among others. These systems deliberately rely on early-layer or fixed representations because deeper contextualized embeddings tend to become increasingly task-specific, thereby distorting global semantic structure and degrading their overall retrieval quality.

To evaluate whether deeper-layer retrieval is beneficial, we incorporated the ranker at *every* layer such that it operates on contextualized embeddings. This intervention substantially degraded performance, causing an average drop of 27.28% across benchmarks, and slowed efficiency by 5.9 \times . These findings reinforce that deeper-layer retrieval is not only computationally prohibitive (as it would require recomputing MaxSim and reassembling contextualized blocks at each layer) but also detrimental for effectiveness.

To summarize, these ablations collectively validate Avey-B’s core architectural principles, namely, (1) decoupling and normalization are critical for both effectiveness and stability (more on this in Appendix J); (2) the residual connection preserves essential local information; (3) the neural compressor delivers substantial efficiency gains with minimal accuracy loss; and (4) the ranker is indispensable, with shallow-embedding retrieval proving both computationally justified and empirically optimal.

I LATENCY RESULTS

In Section 5.4, we reported *throughput* (tokens/second) for Avey-B, ModernBERT, and NeoBERT. In this section, we present *latency* (seconds per forward pass) for the same encoders. As noted in Section 5.4, Avey is recent and lacks a fused-kernel (CUDA/Triton) implementation. Accordingly, we measure Avey-B’s latency using both an *eager* PyTorch implementation (Avey-B-*eager*) and `torch.compile` (Avey-B-*torch-compile*). By contrast, both ModernBERT and NeoBERT have optimized implementations using FlashAttention (Dao et al., 2022). We therefore report their latencies *with* and *without* FlashAttention, and denote the resulting variants as *sys-optimized* and *sys-unoptimized*, respectively.

As with throughput, Avey-B achieves consistently lower latency and markedly superior long-context scaling (see Fig. 4 (a)), underscoring its computational efficiency even in the absence of a fused kernel. To quantify long-context behavior, we again fit a power-law model, $L(N) \propto N^\beta$, where N denotes the sequence length and larger exponents β indicate worse scaling. Under the optimized setting, ModernBERT-sys-optimized and NeoBERT-sys-optimized exhibit $\beta_{\text{ModernBERT}} = 1.17$ and $\beta_{\text{NeoBERT}} = 1.20$, reflecting the bandwidth and memory-pressure limitations inherent to quadratic attention. In contrast, Avey-B-torch-compile attains a substantially milder exponent of $\beta_{\text{Avey-B}} = 0.68$, and delivers more than a 3 \times and 10 \times latency advantage over ModernBERT-sys-optimized and NeoBERT-sys-optimized, respectively, at 96k tokens.

The unoptimized setting further accentuates these differences (see Fig. 4 (b)). ModernBERT-sys-unoptimized and NeoBERT-sys-unoptimized exhibit substantially steeper latency growth, with exponents $\beta_{\text{ModernBERT}} = 1.42$ and $\beta_{\text{NeoBERT}} = 1.63$, and both models encounter out-of-memory failures well before the maximum tested sequence length. In contrast, Avey-B-eager achieves the shallowest growth of all configurations, with $\beta_{\text{Avey-B}} = 0.58$, and maintains stable latency across the entire sequence-length range. These results confirm that Avey-B’s latency advantage is structural, especially since its neural processor depends on split size rather than global sequence length, yielding linear $\mathcal{O}(N)$ scaling and robust long-context performance even in the absence of compiler- or kernel-level optimizations.

J CROSS-SEED VARIANCE ANALYSIS

Table 11: Standard deviations across 10 random seeds for all evaluated encoders and benchmarks.

Model	SC			TC			QA			IR			
	MNLI	QQP	SST-2	CONLL	Onto.	UNER	ReCoRD	SQuAD	SQuADv2	MLDR	MSMARCO	NQ	
Base	Avey-B	0.92	0.12	0.97	0.71	0.12	2.65	0.67	0.17	0.47	0.67	1.34	0.75
	BERT	0.17	0.16	0.31	1.10	0.31	0.69	3.70	0.53	0.32	1.22	0.73	1.44
	RoBERTa	0.13	0.10	0.38	0.25	0.14	0.60	0.42	0.11	0.19	0.32	1.02	0.72
	ModernBERT	0.37	0.12	0.53	0.24	0.11	0.45	0.70	2.36	0.29	1.40	1.39	2.18
M	NeoBERT	0.40	0.14	1.20	0.24	0.17	0.47	5.98	0.91	0.66	4.70	4.30	9.48
	Avey-B	0.20	0.43	0.52	0.33	0.10	1.06	0.27	0.15	0.30	1.03	1.77	1.06
	BERT	0.28	8.24	0.98	0.37	0.13	0.80	2.47	0.97	0.58	0.94	1.54	1.01
	RoBERTa	0.16	0.20	0.41	0.26	0.10	0.61	0.26	0.33	0.18	0.50	1.58	0.74
Large	ModernBERT	0.18	0.08	0.35	0.50	0.14	3.67	17.79	0.25	1.35	2.27	1.70	3.06

To complement the median results reported in Table 2, we further evaluate each model’s robustness by examining its sensitivity to random initialization. As described in Section 5.1, for every benchmark in the SC, TC, QA, and IR categories, we swept four learning rates and fine-tuned each configuration using 10 independent random seeds. While Table 2 reports the *median* performance across seeds at the *best* learning rate for each benchmark, we additionally compute the standard deviation (SD) across the 10 runs for each model–benchmark pair. These SD values quantify the variability induced by initialization and provide an additional perspective on optimization stability and robustness beyond median performance.

At the *base* scale (and *medium* for NeoBERT), RoBERTa exhibits the lowest overall variance, consistent with its well-established fine-tuning stability. Avey-B ranks second, followed by ModernBERT, BERT, and NeoBERT, in that order. Across most benchmarks, Avey-B maintains tightly concentrated variances, with the exception of UNER, where a single outlier resulted in higher variability (SD = 2.65).

At the *large* scale, the differences in stability become more pronounced. ModernBERT, despite strong median performance, exhibits substantial instability on several benchmarks (most notably ReCoRD, UNER, and NQ), suggesting high sensitivity to initialization arising from its alternating attention pattern and extended context window. Likewise, BERT demonstrates occasional catastrophic variance spikes (e.g., QQP with SD = 8.24), indicating susceptibility to poor optimization minima. In contrast, Avey-B maintains uniformly low SDs across nearly all benchmarks, with no signs of pathological instability. Its variances remain tightly bounded (typically below 1.06), often surpassing most Transformer-based baselines and again ranking just behind RoBERTa.

These variance measurements show that Avey-B is among the most statistically consistent encoders in our evaluation. We attribute this robustness to three core architectural principles: (1) the decoupling of static and dynamic layers, which prevents destructive interactions between fixed parameters and similarity scores; (2) row-normalized similarity matrices, which stabilize activation magnitudes and ensure well-behaved gradient flow; and (3) neural compression, which filters out irrelevant signals in retrieved contexts. Collectively, these mechanisms reduce sensitivity to initialization and foster smoother optimization dynamics, accounting for Avey-B’s consistently low variance across tasks.

1296 Table 12: Needle-in-a-haystack (NIAH-1) accuracy across sequence lengths from 1k to 96k for
 1297 several encoders at different scales (M = Medium; OOM = Out-of-Memory).

Model		NIAH-1							
		1k	2k	4k	8k	16k	32k	64k	96k
Base	Avey-B	79.41	79.21	78.94	79.19	78.91	77.73	77.18	75.72
	ModernBERT	67.74	67.64	68.31	70.67	—	—	—	—
M	NeoBERT	79.65	79.13	74.73	—	—	—	—	—
	Avey-B	79.69	79.24	79.03	79.58	79.44	78.44	76.76	76.06
Large	ModernBERT	68.80	67.52	67.20	OOM	—	—	—	—

1308 Table 13: Needle-in-a-haystack (NIAH-2) accuracy across sequence lengths from 1k to 96k for
 1309 several encoders at different scales (M = Medium; OOM = Out-of-Memory).

Model		NIAH-2							
		1k	2k	4k	8k	16k	32k	64k	96k
Base	Avey-B	78.29	79.40	79.77	78.53	78.70	75.50	73.78	71.86
	ModernBERT	66.99	67.25	69.49	70.48	—	—	—	—
M	NeoBERT	79.61	79.52	80.07	—	—	—	—	—
	Avey-B	78.94	79.48	79.99	78.71	79.07	78.31	74.47	74.54
Large	ModernBERT	66.96	67.29	68.07	OOM	—	—	—	—

K LONG-RANGE BENCHMARK RESULTS

1323 In this section, we evaluate the long-context capabilities of Avey-B, ModernBERT, and NeoBERT
 1324 using a synthetic needle-in-a-haystack (NIAH) benchmark formulated as an extractive ques-
 1325 tion-answering (QA) task augmented with position-sensitive reasoning. Each instance consists of a
 1326 passage of a specified length (e.g., 96k tokens) filled with random distractor tokens and *one or more*
 1327 key-value pairs, where the value constitutes the “needle.” The query contains only the key, and the
 1328 model must extract the corresponding value from the passage.

1329 In the single-needle setting, the task measures a model’s ability to semantically locate the correct
 1330 span within an extremely long sequence. In the multi-needle setting, all key-value pairs share
 1331 the same key, and the query explicitly requests the n^{th} occurrence. This removes any semantic
 1332 disambiguation and introduces a position-sensitive reasoning requirement, where the model must
 1333 identify all candidate spans and reason over their order to select the correct needle.

1334 The evaluation set comprises 40% single-needle (pure long-context QA) and 60% two-needle
 1335 (position-sensitive reasoning) examples, thereby jointly assessing semantic retrieval and positional
 1336 reasoning under extreme sequence lengths, with the latter emphasized by design. We further
 1337 construct two variants of the benchmark with different classes of randomly generated needles, namely,
 1338 alphanumeric (NIAH-1) and numeric (NIAH-2), following the setup introduced in Avey (Hammoud
 1339 & Acharya, 2025).

1340 Tables 12 and 13 illustrate results for Avey-B *base* and *large*, ModernBERT *base* and *large*, and
 1341 NeoBERT *medium* (the only publicly available size) on NIAH-1 and NIAH-2 across sequence
 1342 lengths from 1k to 96k tokens. For ModernBERT and NeoBERT, we show results only up to their
 1343 respective trained context windows (i.e., 8k and 4k tokens). In contrast, Avey-B imposes no fixed
 1344 maximum sequence length and generalizes seamlessly beyond its trained 2,048-token context
 1345 window, enabling evaluation up to 96k tokens.

1346 On NIAH-1 (Table 12), Avey-B exhibits strong robustness and scalability relative to ModernBERT
 1347 and NeoBERT. Both Avey-B *base* and *large* maintain near-constant accuracy from 1k to 96k to-
 1348 kens, with only a modest 3–4 point decrease over a 96 \times increase in sequence length. This stability
 1349 demonstrates that Avey-B effectively resolves long-range dependencies and generalizes far beyond
 its trained context window. By comparison, ModernBERT and NeoBERT cannot operate beyond

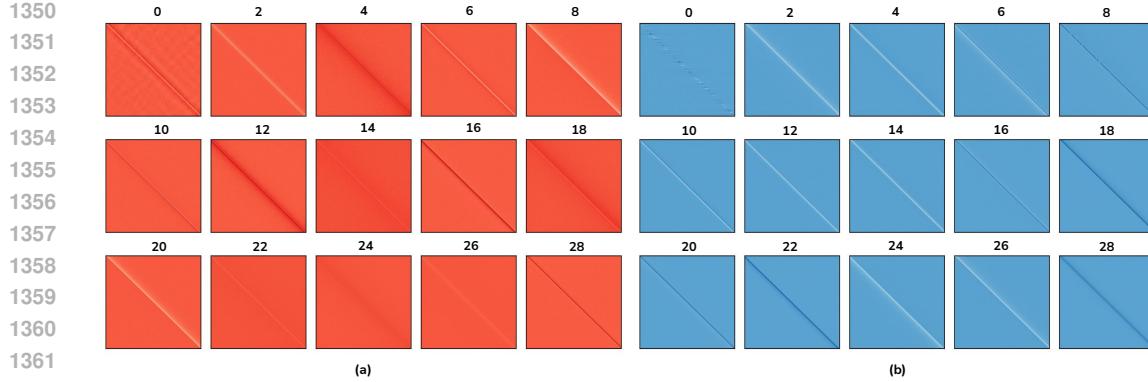


Figure 5: Learned static cross-embedding projection matrices for the (a) *coupled* configuration (left or red) with 15 matrices uniformly subsampled from 30 static layers and (b) *decoupled* configuration (right or blue) with all 15 static matrices (dynamic and static layers are interleaved, hence, only 15 static matrices exist). For comparability, we display 15 layers per panel. The coupled setting exhibits diffuse, more homogeneous patterns (e.g., see layers 14, 22, 24, and 26) suggestive of redundancy, whereas the decoupled setting shows sharper, more heterogeneous structure and variability in spread, indicating greater representational diversity.

their 8k-token and 4k-token trained windows, respectively. NeoBERT matches Avey-B at very short contexts (1–2k) but drops by roughly 5 points at 4k, while ModernBERT lags behind Avey-B by 10–12 points even at short sequence lengths². Moreover, ModernBERT *large* fails at 8k tokens due to out-of-memory issues even on NVIDIA B200 GPUs using the smallest feasible batch size.

Results on NIAH-2 (Table 13) mirror the trends observed on NIAH-1 and further reinforce Avey-B’s long-context robustness. Both Avey-B *base* and *large* maintain strong performance across extreme sequence lengths, with Avey-B *base* decreasing from 78.3 at 1k to 71.9 at 96k tokens, and Avey-B *large* remaining similarly stable (78.9 → 74.5). ModernBERT, by contrast, trails Avey-B by 9–12 points at short contexts and fails at 8k tokens due to memory limitations, preventing any long-context evaluation. NeoBERT remains competitive at very short lengths (1–4k) but, like ModernBERT, cannot operate beyond its 4k-token window, and therefore cannot be evaluated in the long-context regime.

In summary, Avey-B is the *only* model capable of sustaining high accuracy up to 96k tokens on this question-answering, reasoning-intensive benchmark, despite being trained with a context window of only 2,048 tokens. These results demonstrate that Avey-B not only generalizes far beyond its trained context width, but also preserves long-range reasoning fidelity in regimes where existing Transformer-based encoders either fail to extrapolate or collapse entirely.

L COUPLED VS. DECOUPLED LAYERS: A STATISTICAL ANALYSIS

We now analyze the learned cross-embedding projection matrices (say, \mathbf{V}) for the *coupled* and *de-coupled* Avey-B models from the ablation study in Appendix H. Table 14 reports summary statistics for each model. In the coupled case, we observe a clear *positivity bias*, especially in deeper layers, wherein the fraction of positive entries (i.e., the number of positive weights divided by the total number of weights) approaches one in several layers (in layers 8 and 13, it indeed hit 1). This bias can be explained as follows. Because the enricher employs a nonnegative activation (i.e., ReLU²), the contextualizer’s similarity matrix (say, \mathbf{S}) is elementwise nonnegative. As such, the coupled mixing $\mathbf{M} = \mathbf{V} \odot \mathbf{S}$ inherits its *signs* entirely from \mathbf{V} . Any negative entry in \mathbf{V} flips a large positive similarity into a negative contribution, violating *monotonicity with respect to relevance* (see Appendix A) and degrading training. The optimizer therefore pushes \mathbf{V} toward nonnegativity to avoid these destructive sign inversions, yielding the observed late-layer collapse toward positive weights.

²For ModernBERT versus NeoBERT, we hypothesize that the consistently lower scores of ModernBERT stem from its local–global alternating attention pattern, compared to the full bidirectional self-attention used in NeoBERT.

1404
 1405 Despite this positivity bias, a nontrivial fraction of negative entries persists in the coupled model
 1406 (see Table 14 again). This residual negativity is precisely the failure mode our hypothesis predicts,
 1407 that is, wherever a neuron retains negative weights, large positive similarities can be inverted into
 1408 negative contributions, yielding local violations of relevance monotonicity.

1409 By contrast, in the decoupled case the dynamic layers alone produce the mixing weights. These
 1410 weights are normalized and nonnegative by construction, so monotonicity is enforced at the similarity
 1411 operator. The static layers are learned separately and no longer need to be driven into nonnegativity
 1412 to preserve monotonicity. As shown in Table 14, this yields a near-zero mean with roughly
 1413 balanced positive and negative weights (without the late-layer positivity bias), retains inhibitory
 1414 patterns (i.e., learned negative influences) where useful, and avoids the sign-flip failure mode.

1415 Beyond sign distribution, the two models also diverge in the *dispersion* of their weights. Coupled
 1416 matrices exhibit reduced standard deviation across layers, indicating more stable transformations
 1417 that converge toward smooth and homogeneous patterns. Decoupled matrices, by contrast, sus-
 1418 tain larger fluctuations, admitting both stronger positive and stronger negative values. This higher
 1419 variance may reflect greater representational flexibility. Norm statistics supports this interpreta-
 1420 tion, whereby coupled matrices accumulate larger ℓ_1 norm, distributing weight more evenly across
 1421 entries, whereas decoupled matrices attain slightly higher ℓ_2 values, implying that fewer entries
 1422 dominate with sharper magnitudes.

1423 Qualitatively, the static matrices in both variants exhibit Toeplitz-like (approximately shift-invariant)
 1424 structure reminiscent of gMLP (Liu et al., 2021) (see Fig. 5). As in gMLP, where such patterns
 1425 emerge without an explicit prior, our static layers converge to diagonally dominant, near-diagonal
 1426 matrices indicative of locality. This alignment suggests that locality-preserving, Toeplitz-like struc-
 1427 ture can arise naturally in architectures that employ fixed, input-independent transformations to
 1428 stabilize and scaffold subsequent dynamic computations.

1429 In summary, *coupling* tends to regularize the cross-embedding projections toward homogeneous,
 1430 nearly nonnegative transformations, whereas *decoupling* promotes healthy diversity and sharper
 1431 structure, while preserving monotonicity with respect to relevance.

1432

1433

1434

1435

1436

1437

1438

1439

1440

1441

1442

1443

1444

1445

1446

1447

1448

1449

1450

1451

1452

1453

1454

1455

1456

1457

1458
1459
1460
1461
1462
14631464 Table 14: Layer statistics for coupled vs. decoupled settings. For comparability, we display 15 layers
1465 per setting. For the coupled setting, we uniformly subsampled 15 layers from 30 static layers. For
1466 the decoupled setting, *all* the 15 static layers are shown (dynamic and static layers are interleaved,
1467 hence, only 15 static layers exist). The coupled setting exhibits *positivity bias* (see the “fraction of
1468 positive” column), while the decoupled setting demonstrates more *balanced* positive and negative
1469 weights, indicating greater representational diversity.1470
1471
1472
1473
1474
1475
1476
1477
1478
1479
1480
1481
1482
1483
1484
1485
1486
1487
1488

Layer	Coupled									
	Mean	Std	Min	Median	Max	Abs. Mean	L1 Norm	L2 Norm	Frac. Pos.	Frac. Neg.
1	0.00	0.12	-1.67	0.00	1.20	0.05	3369.19	30.68	0.47	0.53
2	-0.01	0.11	-1.55	0.00	0.30	0.04	2345.75	27.73	0.53	0.47
3	0.08	0.09	-0.23	0.06	0.68	0.09	5906.35	31.56	0.90	0.10
4	0.00	0.10	-1.19	0.01	0.29	0.03	2211.32	24.42	0.61	0.39
5	-0.03	0.13	-1.09	0.00	0.15	0.06	3724.35	33.82	0.55	0.45
6	0.01	0.08	-1.47	0.00	0.63	0.03	1963.87	21.33	0.57	0.43
7	0.03	0.11	-0.29	0.00	1.00	0.06	3959.52	30.11	0.52	0.49
8	0.11	0.04	-0.14	0.11	0.33	0.11	7525.47	31.17	1.00	0.00
9	0.00	0.10	-0.38	-0.01	1.43	0.04	2300.72	25.97	0.34	0.66
10	0.09	0.06	-0.07	0.08	0.38	0.09	5844.50	27.45	0.98	0.02
11	-0.02	0.10	-0.90	0.01	0.12	0.04	2604.39	26.53	0.59	0.41
12	0.07	0.02	-0.04	0.07	0.16	0.07	4697.84	19.12	0.98	0.02
13	0.05	0.03	-0.04	0.05	0.21	0.05	3487.26	15.28	1.00	0.00
14	0.03	0.02	-0.14	0.03	0.18	0.04	2293.24	9.56	0.95	0.05
15	0.00	0.07	-0.39	-0.01	0.83	0.03	1912.78	18.29	0.34	0.66
Avg.	0.03	0.08	-0.64	0.03	0.53	0.06	3609.77	24.87	0.69	0.31

1489
1490
1491
1492
1493
1494
1495
1496
1497
1498
1499
1500
1501
1502
1503
1504
1505
1506

Layer	Decoupled									
	Mean	Std	Min	Median	Max	Abs. Mean	L1 Norm	L2 Norm	Frac. Pos.	Frac. Neg.
1	0.00	0.08	-0.98	0.00	1.04	0.04	2327.62	21.31	0.51	0.49
2	-0.02	0.12	-1.27	0.00	0.23	0.05	3065.94	30.83	0.51	0.49
3	-0.01	0.12	-1.93	0.01	0.27	0.04	2312.23	30.24	0.66	0.34
4	-0.01	0.12	-1.01	0.01	1.35	0.04	2825.07	29.60	0.60	0.40
5	0.00	0.10	-0.50	-0.01	1.70	0.03	2129.54	26.59	0.42	0.58
6	0.00	0.11	-2.09	0.00	0.68	0.03	1978.08	28.17	0.61	0.39
7	0.00	0.11	-1.60	0.01	0.33	0.03	2050.08	27.01	0.69	0.31
8	-0.01	0.11	-1.22	0.01	0.30	0.03	2104.69	27.52	0.70	0.30
9	0.00	0.09	-0.26	0.00	1.37	0.03	1697.87	23.67	0.42	0.58
10	0.01	0.12	-0.34	-0.01	1.27	0.03	2145.28	30.16	0.34	0.66
11	0.00	0.09	-0.28	-0.01	1.18	0.03	1775.56	23.28	0.35	0.65
12	0.01	0.11	-0.22	-0.01	1.62	0.03	2019.32	29.24	0.33	0.67
13	-0.02	0.11	-1.35	0.00	0.24	0.04	2637.36	28.99	0.58	0.42
14	-0.02	0.10	-1.04	0.00	0.18	0.04	2292.65	26.27	0.61	0.39
15	0.01	0.08	-0.06	0.00	0.82	0.03	1779.38	19.60	0.51	0.49
Avg.	0.00	0.10	-0.94	0.00	0.84	0.03	2209.38	26.83	0.52	0.48

1507
1508
1509
1510
1511