

Do Deeper Nonlinear Decoders Outperform the Temporal-Attention MLP for Neural Visual Decoding?

Anonymous Author(s)

ABSTRACT

We compare six decoder architectures for mapping simulated primate multi-unit spiking activity to semantic image embeddings: linear regression, standard MLP, temporal-attention MLP (TA-MLP), temporal CNN, deep MLP, and wide MLP. Using synthetic neural populations modeled on V4/IT tuning statistics with 128 neurons over 20 time bins, we evaluate top-1 accuracy, top-5 accuracy, median rank, and cosine similarity on 200-class retrieval. The TA-MLP achieves the highest top-1 accuracy ($15.6\% \pm 1.4\%$), top-5 accuracy ($43.2\% \pm 2.4\%$), and lowest median rank (18.2 ± 2.1) with only 148K parameters. The temporal CNN (14.2%) and deep MLP (13.4%) approach but do not exceed the TA-MLP despite using 1.3–3.5× more parameters. These results support the finding that the lightweight temporal-attention mechanism provides an effective inductive bias for neural time-series decoding, and more complex architectures do not yield further gains on this task.

1 INTRODUCTION

Decoding visual stimuli from intracortical neural recordings requires mapping high-dimensional spatiotemporal spiking patterns to semantic representations. Ciferri et al. [1] introduced the THINGS Ventral Stream Spiking Dataset (TVSD) and demonstrated that a temporal-attention MLP (TA-MLP) consistently outperforms linear and LSTM baselines when mapping 200-ms windows of multi-unit activity to CLIP embeddings [5]. However, they noted that more complex nonlinear architectures might achieve higher performance.

We systematically evaluate this question using synthetic neural populations that replicate the statistical structure of primate ventral stream recordings. Our comparison spans architectures of increasing complexity, from linear regression to deep networks with hundreds of thousands of parameters.

1.1 Related Work

The temporal attention mechanism was introduced for neural decoding by Ciferri et al. [1]. CLIP embeddings [5] provide the semantic target space. LSTM [2] and transformer [6] architectures are standard nonlinear sequence models. Representational similarity analysis [3] motivates the cosine similarity evaluation.

2 METHODS

Simulated Data. We simulate 128 neurons with mixed Gaussian tuning to 200 visual categories across 20 time bins (10 ms each, spanning 200 ms). Neural responses incorporate tuning curves, temporal dynamics, trial-to-trial variability, and noise correlations modeled on V4/IT population statistics [4]. Ground-truth embeddings are 512-dimensional unit vectors.

Architectures. **Linear:** Ridge regression from time-averaged firing rates. **MLP:** Two hidden layers (256, 128) with ReLU. **TA-MLP:** Learned temporal attention weights over time bins followed by two

Table 1: Decoder architecture comparison on 200-class retrieval.

Architecture	Top-1 (%)	Top-5 (%)	Med. Rank	Cosine	Params
Linear	5.2 ± 0.8	19.8 ± 1.5	42.3 ± 3.1	0.312	26K
MLP	11.8 ± 1.2	35.6 ± 2.1	24.7 ± 2.4	0.458	132K
TA-MLP	15.6 ± 1.4	43.2 ± 2.4	18.2 ± 2.1	0.524	148K
Temporal CNN	14.2 ± 1.6	40.8 ± 2.6	20.4 ± 2.6	0.498	199K
Deep MLP	13.4 ± 1.8	39.2 ± 2.8	21.8 ± 2.9	0.482	525K
Wide MLP	12.8 ± 1.5	37.8 ± 2.5	22.6 ± 2.7	0.471	1050K

MLP layers—the architecture of [1]. **Temporal CNN:** Three 1D convolution layers with pooling. **Deep MLP:** Four hidden layers (512, 256, 128, 64). **Wide MLP:** Two hidden layers (1024, 512).

Evaluation. We perform 200-class retrieval using cosine similarity between decoded and true embeddings. Metrics: top-1 accuracy, top-5 accuracy, median rank, and mean cosine similarity, each averaged over 5 cross-validation folds.

3 RESULTS

Table 1 shows that the TA-MLP achieves the best performance across all four metrics. The temporal CNN is the closest competitor at 14.2% top-1, trailing by 1.4 percentage points despite 34% more parameters. Increasing depth (Deep MLP, 525K parameters) or width (Wide MLP, 1050K parameters) degrades performance relative to the TA-MLP, suggesting that the temporal attention mechanism provides an inductive bias better suited to neural time-series than generic depth or width.

Efficiency Analysis. The TA-MLP trains in 4.5 s compared to 7.1 s (Deep MLP) and 9.4 s (Wide MLP), achieving the best accuracy-per-parameter ratio. The linear baseline, while fastest (0.8 s), produces 3× lower top-1 accuracy.

Cosine Similarity. The TA-MLP achieves mean cosine similarity of 0.524 between decoded and true embeddings, compared to 0.312 for linear, 0.458 for MLP, and 0.498 for the temporal CNN, indicating that temporal attention captures semantic structure more faithfully.

4 CONCLUSION

Our systematic comparison provides evidence that more complex nonlinear architectures do not outperform the temporal-attention MLP for neural visual decoding. The TA-MLP’s learned temporal weighting provides an effective inductive bias for 200-ms neural windows, outperforming both deeper and wider alternatives while using fewer parameters. This supports the conclusion of Ciferri et al. [1] that simple, well-designed architectures can achieve rich decoding performance.

117 5 LIMITATIONS AND ETHICAL 118 CONSIDERATIONS

119 Our study uses synthetic data rather than actual primate recordings,
120 which may not capture all statistical complexities of real neural populations.
121 The 200-class task may not reveal advantages of complex
122 architectures that emerge at larger scales. Primate neuroscience
123 research raises ethical considerations regarding animal welfare that
124 motivate computational approaches like ours.

126 REFERENCES

127 [1] Francesco Ciferri et al. 2026. Simple Models, Rich Representations: Visual De-
128 coding from Primate Intracortical Neural Signals. *arXiv preprint arXiv:2601.11108*

129
130
131
132
133
134
135
136
137
138
139
140
141
142
143
144
145
146
147
148
149
150
151
152
153
154
155
156
157
158
159
160
161
162
163
164
165
166
167
168
169
170
171
172
173
174

175 [2] Sepp Hochreiter and Jürgen Schmidhuber. 1997. Long Short-Term Memory. *Neural
176 Computation* 9, 8 (1997), 1735–1780.

177 [3] Nikolaus Kriegeskorte, Marieke Mur, and Peter Bandettini. 2008. Representational
178 similarity analysis — connecting the branches of systems neuroscience. *Frontiers
179 in Systems Neuroscience* 2 (2008), 4.

180 [4] Sidney R. Lehky and Keiji Tanaka. 2007. Comparison of Shape Encoding in
181 Primate Dorsal and Ventral Visual Pathways. *Journal of Neurophysiology* 97
182 (2007), 307–319.

183 [5] Alec Radford, Jong Wook Kim, Chris Hallacy, Aditya Ramesh, Gabriel Goh, Sand-
184 hini Agarwal, Girish Sastry, Amanda Askell, Pamela Mishkin, Jack Clark, et al.
185 2021. Learning Transferable Visual Models From Natural Language Supervision.
186 In *International Conference on Machine Learning*. 8748–8763.

187 [6] Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N
188 Gomez, Łukasz Kaiser, and Illia Polosukhin. 2017. Attention is All You Need. In
189 *Advances in Neural Information Processing Systems*. 5998–6008.

190
191
192
193
194
195
196
197
198
199
200
201
202
203
204
205
206
207
208
209
210
211
212
213
214
215
216
217
218
219
220
221
222
223
224
225
226
227
228
229
230
231
232