

# FEATURES EMERGE AS DISCRETE STATES: THE FIRST APPLICATION OF SAEs TO 3D REPRESENTATIONS

000  
001  
002  
003  
004  
005 **Anonymous authors**  
006 Paper under double-blind review  
007  
008  
009  
010

## ABSTRACT

011 Sparse Autoencoders (SAEs) are a powerful dictionary learning technique for de-  
012 composing neural network activations, translating the hidden state into human  
013 ideas with high semantic value despite no external intervention or guidance. How-  
014 ever, this technique has rarely been applied outside of the textual domain, limiting  
015 theoretical explorations of feature decomposition. We present the **first applica-**  
016 **tion of SAEs to the 3D domain**, analyzing the features used by a state-of-the-art  
017 3D reconstruction VAE applied to 53k 3D models from the Objaverse dataset. We  
018 observe that the network encodes discrete rather than continuous features, leading  
019 to our key finding: **such models approximate a discrete state space, driven by**  
020 **phase-like transitions from feature activations**. Through this state transition  
021 framework, we address three otherwise unintuitive behaviors — the inclination  
022 of the reconstruction model towards positional encoding representations, the sig-  
023 moidal behavior of reconstruction loss from feature ablation, and the bimodality  
024 in the distribution of phase transition points. This final observation suggests the  
025 **model redistributes the interference caused by superposition to prioritize the**  
026 **saliency of different features**. Our work not only compiles and explains unex-  
027 pected phenomena regarding feature decomposition, but also provides a frame-  
028 work to explain the model’s feature learning dynamics. The code and dataset of  
029 encoded 3D objects will be available on release.  
030

## 1 INTRODUCTION

031 Interpretability research has recently focused on translating a model’s latent state into sets of human-  
032 readable concepts. To this end, studies have used sparse autoencoders (SAEs) as a dictionary-  
033 learning tool applied to the latent vectors of LLMs (Bricken et al. 2023). These publications find se-  
034 mantically interpretable pipelines in foundational models for a myriad of tasks, including arithmetic  
035 (Lindsey et al. 2025), protein characteristics (Garcia & Ansuini 2025), and image-text relationships  
036 (Yan et al. 2025).  
037

038 The success of these methods suggests that the hidden state functions as a concept space, where  
039 individual axes correspond to independent features, allowing the model to generalize (Elhage et al.  
040 2023). Furthermore, models can layer a number of feature vectors far greater than the cardinality of  
041 the latent space through a process called superposition, at the cost of interference from compression  
042 (Hänni et al. 2024). However, research in feature decomposition is lacking in two key areas. First,  
043 the scope of data domains has been limited — **recent feature decomposition techniques, particularly**  
044 **SAEs, have rarely been applied to industries that use unordered data with continuous features, would**  
045 **benefit from improved transparency**. Second, a model’s learned features are often counter-intuitively  
046 constructed — existing research tends to empirically discuss *what* features contribute to a model’s  
047 performance, rather than explore *why* or *how* these features were chosen by the model. **We believe**  
048 **this problem is exacerbated by the focus on textual data, which draws input data from a finite,**  
049 **discrete vocabulary**. We cite further studies in Appendix A.

050 Our work seeks to address these gaps. For the first gap, we are, to the best of our knowledge, the  
051 first work to apply an SAE on latent vectors handling 3D data; specifically, 3D models sampled from  
052 Objaverse (Deitke et al. 2022) encoded with Dora-VAE (Chen et al. 2024). **3D data is a domain well-**  
053 **suited for feature decomposition research, because a) it is visually obvious when detected features**  
054 **have semantic meaning, b) existing datasets have a wide variety of immediately recognizable objects**

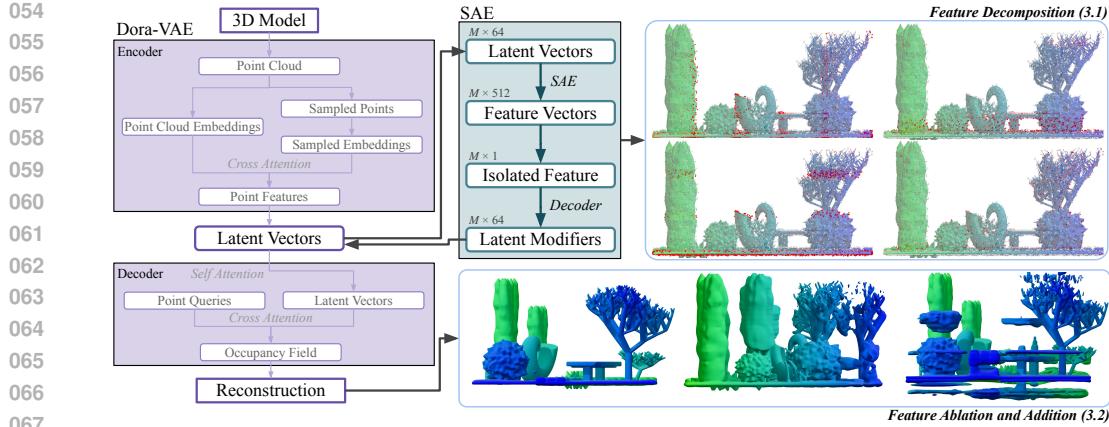


Figure 1: **Our feature decomposition pipeline.** Dora-VAE is a 3D reconstruction model, encoding 3D objects to  $M$  latent vectors each. We apply our SAE to these latent vectors, decomposing each vector to a linear combination of features. We can visualize the effects of an individual feature by plotting the its presence in each latent vector (Section 3.1) or modifying the latent vector and observing the effects on the reconstruction (Section 3.2).

with unique semantic combinations c) many large industries (animation, design, architecture, etc.) rely on AI tools in this domain, and d) we avoid datasets that are noisy, synthetic, homogenous, sparse, and/or constructed for toy experiments. In addition, 3D data draws from an unordered, continuous domain, i.e. the position of points in 3D space. This is a qualitatively distinct challenge to text, which results in interesting feature dynamics.

We find clear, human-interpretable features in the latent vectors and report on their semantic meaning. We also observe that Dora-VAE learns feature representations that we would consider unorthodox. Namely, positional information is represented *discretely rather than continuously*, transition points between high-impact states follow a *unimodal distribution*, and transition points between low-impact states follow a *bimodal distribution*.

These idiosyncrasies present an opportunity to address the second gap in research. Where previous work prioritized the identification of a model’s conceptual pipeline, we study the **learning dynamics** behind this pipeline. We deconstruct the optimization step and identify two terms that independently attend to the **presence** and **identity** of individual features. The dichotomy between these terms offers explanations to the learning behaviors of Dora-VAE highlighted earlier — particularly, we suggest the unimodality of high-impact transition states is explained through the presence term, and the bimodality of low-impact transition states is explained through the redistribution of superposition interference. This framework is potentially universally applicable, intended to provide context for future interpretability work to discuss how concepts form a discrete state space.

We substantiate our framework through a set of verifying experiments. First, we establish the veracity of our learned features by highlighting noteworthy features and demonstrating the effects of targeted feature intervention. We show the clear semantic effect a feature has on a reconstructed output. Second, we observe and discuss several counter-intuitive behaviors that are explained by our framework. We do so through a series of 848k independent feature interventions across a set of 53k 3D models to observe patterns of changes in loss.

Ultimately, we provide the first application of SAEs to 3D data and explain the unusual properties of the feature space exhibited through a novel theoretical framework. In future work, we hope to evaluate the universality and consistency of our framework on different models and modalities, as well as further investigate the feature dynamics of models in the 3D domain.

108  
109  

## 2 FEATURE DECOMPOSITION PRELIMINARIES

110  
111  
112  
We describe a model as the composition of functions  $f : \mathbf{x} \mapsto \mathbf{z}$  and  $g : \mathbf{z} \mapsto \mathbf{y}$ , where  $\mathbf{z}$  is a latent  
vector. Given a dataset  $\{(\mathbf{x}_i, \mathbf{y}_i)\}_{i=1}^N$ , the training objective is to optimize parameters  $\theta_f$  and  $\theta_g$  by  
minimizing a loss function  $\mathcal{L}$  over the dataset:

113  
114  
115  
$$\min_{\theta_f, \theta_g} \sum_{i=1}^N \mathcal{L}(g(\theta_g; f(\theta_f; \mathbf{x}_i)), \mathbf{y}_i) \quad (1)$$

116  
117  
118  
Several works have suggested theoretically (Bengio et al. 2013) and empirically (Elhage et al.  
119  
2022) that latent representations in both humans and models can be viewed sets of semantic ideas  
120  
— borrowing terminology proposed by Kim et al. (2018), we denote these ideas as a set of vectors  
121  
 $\mathbf{E}$ , and the space spanned by these vectors as  $E$ . These works suggest that  $\mathbf{z} \in E$  and that  $\mathbf{z}$  can be  
122  
decomposed to a linear combination of vectors from  $\mathbf{E}$ :

123  
124  
$$\mathbf{z} = f(\theta_f; \mathbf{x}) = \mathbf{E}(\theta_f)^T \boldsymbol{\alpha}(\theta_f; \mathbf{x}) \quad \text{where } \mathbf{E} = [\mathbf{e}_1, \dots, \mathbf{e}_n] \subset E, \quad \boldsymbol{\alpha} = (\alpha_1, \dots, \alpha_n)^T \quad (2)$$

125  
126  
Abstractly, we state that a model’s latent space interprets a given input  $\mathbf{x}$  as a set of scalars  $\boldsymbol{\alpha}$   
127  
modifying a set of learned features  $\mathbf{E}$ . We refer to  $\alpha_j$  as the *presence* of feature  $j$  and  $\mathbf{e}_j$  as the  
128  
*identity* of feature  $j$ .129  
130  
Even if an input  $\mathbf{x}$  is out of the domain of the training dataset, the model still attempts, and likely  
131  
fails, to frame the input in these features. For example, image adversarial attacks use noise that are  
132  
completely out-of-distribution, but a classification model must still estimate the presence of each  
133  
feature  $\mathbf{e}_j$  (Gorton & Lewis 2025). If the resulting feature presences is similar to an in-distribution  
134  
input, the image is misclassified.135  
136  
Recent LLM studies (Bricken et al. 2023) use a sparse autoencoder (SAE) to approximate this  
137  
decomposition with the assumption that  $\boldsymbol{\alpha}$  is sparse; that is, the number of feature vectors in  $\mathbf{E}$  (i.e.  
138  
the dictionary size) is large compared to the number of features needed to represent a vector in  $E$ .  
139  
Given a collection of input  $\mathbf{x}$  and their corresponding latent vectors  $\mathbf{z}$ , we attempt to approximate  $\boldsymbol{\alpha}$   
140  
and  $\mathbf{E}$  through the following parametrization, known as a BatchTopK SAE (Bussmann et al. 2024):

141  
142  
143  
144  
145  
146  
147  
148  
For fixed  $\theta_f$  :  $\boldsymbol{\alpha}(\theta_f, \mathbf{x}) \approx \text{Enc}(\mathbf{z})$  where  $\text{Enc}(\mathbf{z}) = \text{TopK}(\mathbf{W}^{Enc} \mathbf{z} + \mathbf{b}^{Enc})$   
149  
 $\mathbf{E}(\theta_f) \approx \mathbf{W}^{Dec}$  where  $\hat{\mathbf{z}} = \mathbf{W}^{Dec} \text{Enc}(\mathbf{z}) + \mathbf{b}^{Dec}$  (3)

150  
151  
152  
153  
154  
155  
156  
157  
158  
159  
160  
161  
where TopK selects the top  $k$  nonzero values across the batch. The linear weight matrix  $\mathbf{W}^{Dec}$   
162  
approximates the set of features  $\mathbf{E}$ , forming an overcomplete dictionary. Thus,  $\text{Enc}(\mathbf{z})$  is a sparse  
163  
representation of  $\mathbf{z}$  using  $\mathbf{W}^{Dec}$  as the set of feature vectors. We train with standard reconstruction  
164  
loss, alongside an auxiliary loss based on the reconstruction from dead features —  $\hat{\mathbf{z}}_{dead}$  is the  
165  
reconstruction using only dead features, and  $\beta$  is a scalar hyperparameter (Gao et al. 2025).

166  
167  
168  
169  
170  
$$\begin{aligned} \mathcal{L}(\theta_{SAE}) &= \mathcal{L}_{recon}(\mathbf{z}, \hat{\mathbf{z}}) + \beta \mathcal{L}_{recon}(\mathbf{z}, \hat{\mathbf{z}}_{dead}) \\ \mathcal{L}_{recon}(\mathbf{z}, \hat{\mathbf{z}}) &= \sum_i \|z_i - \hat{z}_i\|_2^2 \end{aligned} \quad (4)$$

171  
172  
173  

### 2.1 THE LEARNING DYNAMICS OF FEATURE DECOMPOSITION

174  
175  
176  
177  
178  
179  
180  
181  
182  
183  
184  
185  
186  
187  
188  
189  
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  
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  
260  
261  
262  
263  
264  
265  
266  
267  
268  
269  
270  
271  
272  
273  
274  
275  
276  
277  
278  
279  
280  
281  
282  
283  
284  
285  
286  
287  
288  
289  
290  
291  
292  
293  
294  
295  
296  
297  
298  
299  
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  
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  
250  
251  
252  
253  
254  
255  
256  
257  
258  
259  
260  
261  
262  
263  
264  
265  
266  
267  
268  
269  
270  
271  
272  
273  
274  
275  
276  
277  
278  
279  
280  
281  
282  
283  
284  
285  
286  
287  
288  
289  
290  
291  
292  
293  
294  
295  
296  
297  
298  
299  
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  
220  
221  
222  
223  
224  
225  
226  
227  
228  
229  
230  
231  
232  
233  
234  
235  
236  
237  
238  
239  
230  
231  
232  
233  
234  
235  
236  
237  
238  
239  
240  
241  
242  
243  
244  
245  
246  
247  
248  
249  
240  
241  
242  
243  
244  
245  
246  
247  
248  
249  
250  
251  
252  
253  
254  
255  
256  
257  
258  
259  
260  
261  
262  
263  
264  
265  
266  
267  
268  
269  
270  
271  
272  
273  
274  
275  
276  
277  
278  
279  
280  
281  
282  
283  
284  
285  
286  
287  
288  
289  
290  
291  
292  
293  
294  
295  
296  
297  
298  
299  
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  
233  
234  
235  
236  
237  
238  
239  
240  
241  
242  
243  
244  
245  
246  
247  
248  
249  
240  
241  
242  
243  
244  
245  
246  
247  
248  
249  
250  
251  
252  
253  
254  
255  
256  
257  
258  
259  
260  
261  
262  
263  
264  
265  
266  
267  
268  
269  
270  
271  
272  
273  
274  
275  
276  
277  
278  
279  
280  
281  
282  
283  
284  
285  
286  
287  
288  
289  
290  
291  
292  
293  
294  
295  
296  
297  
298  
299  
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  
233  
234  
235  
236  
237  
238  
239  
240  
241  
242  
243  
244  
245  
246  
247  
248  
249  
240  
241  
242  
243  
244  
245  
246  
247  
248  
249  
250  
251  
252  
253  
254  
255  
256  
257  
258  
259  
260  
261  
262  
263  
264  
265  
266  
267  
268  
269  
270  
271  
272  
273  
274  
275  
276  
277  
278  
279  
280  
281  
282  
283  
284  
285  
286  
287  
288  
289  
290  
291  
292  
293  
294  
295  
296  
297  
298  
299  
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  
233  
234  
235  
236  
237  
238  
239  
240  
241  
242  
243  
244  
245  
246  
247  
248  
249  
240  
241  
242  
243  
244  
245  
246  
247  
248  
249  
250  
251  
252  
253  
254  
255  
256  
257  
258  
259  
260  
261  
262  
263  
264  
265  
266  
267  
268  
269  
270  
271  
272  
273  
274  
275  
276  
277  
278  
279  
280  
281  
282  
283  
284  
285  
286  
287  
288  
289  
290  
291  
292  
293  
294  
295  
296  
297  
298  
299  
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  
233  
234  
235  
236  
237  
238  
239  
240  
241  
242  
243  
244  
245  
246  
247  
248  
249  
240  
241  
242  
243  
244  
245  
246  
247  
248  
249  
250  
251  
252  
253  
254  
255  
256  
257  
258  
259  
260  
261  
262  
263  
264  
265  
266  
267  
268  
269  
270  
271  
272  
273  
274  
275  
276  
277  
278  
279  
280  
281  
282  
283  
284  
285  
286  
287  
288  
289  
290  
291  
292  
293  
294  
295  
296  
297  
298  
299  
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  
233  
234  
235  
236  
237  
238  
239  
240  
241  
242  
243  
244  
245  
246  
247  
248  
249  
240  
241  
242  
243  
244  
245  
246  
247  
248  
249  
250  
251  
252  
253  
254  
255  
256  
257  
258  
259  
260  
261  
262  
263  
264  
265  
266  
267  
268  
269  
270  
271  
272  
273  
274  
275  
276  
277  
278  
279  
280  
281  
282  
283  
284  
285  
286  
287  
288  
289  
290  
291  
292  
293  
294  
295  
296  
297  
298  
299  
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  
233  
234  
235  
236  
237  
238  
239  
240  
241  
242  
243  
244  
245  
246  
247  
248  
249  
240  
241  
242  
243  
244  
245  
246  
247  
248  
249  
250  
251  
252  
253  
254  
255  
256  
257  
258  
259  
260  
261  
262  
263  
264  
265  
266  
267  
268  
269  
270  
271  
272  
273  
274  
275  
276  
277  
278  
279  
280  
281  
282  
283  
284  
285  
286  
287  
288  
289  
290  
291  
292  
293  
294  
295  
296  
297  
298  
299  
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  
233  
234  
235  
236  
237  
238  
239  
240  
241  
242  
243  
244  
245  
246  
247  
248  
249  
240  
241  
242  
243  
244  
245  
246  
247  
248  
249  
250  
251  
252  
253  
254  
255  
256  
257  
258  
259  
260  
261  
262  
263  
264  
265  
266  
267  
268  
269  
270  
271  
272  
273  
274  
275  
276  
277  
278  
279  
280  
281  
282  
283  
284  
285  
286  
287  
288  
289  
290  
291  
292  
293  
294  
295  
296  
297  
298  
299  
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  
233  
234  
235  
236  
237  
238  
239  
240  
241  
242  
243  
244  
245  
246  
247  
248  
249  
240  
241  
242  
243  
244  
245  
246  
247  
248  
249  
250  
251  
252  
253  
254  
255  
256  
257  
258  
259  
260  
261  
262  
263  
264  
265  
266  
267  
268  
269  
270  
271  
272  
273  
274  
275  
276  
277  
278  
279  
280  
281  
282  
283  
284  
285  
286  
287  
288  
289  
290  
291  
292  
293  
294  
295  
296  
297  
298  
299  
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  
233  
234  
235  
236  
237  
238  
239  
240  
241  
242  
243  
244  
245  
246  
247  
248  
249  
240  
241  
242  
243  
244  
245  
246  
247  
248  
249  
250  
251  
252  
253  
254  
255  
256  
257  
258  
259  
260  
261  
262  
263  
264  
265  
266  
267  
268  
269  
270  
271  
272  
273  
274  
275  
276  
277  
278  
279  
280  
281  
282  
283  
284  
285  
286  
287  
288  
289  
290  
291  
292  
293  
294  
295  
296  
297  
298  
299  
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  
233  
234  
235  
236  
237  
238  
239  
240  
241  
242  
243  
244  
245  
246  
247  
248  
249  
240  
241  
242  
243  
244  
245  
246  
247  
248  
249  
250  
251  
252  
253  
254  
255  
256  
257  
258  
259  
260  
261  
262  
263  
264  
265  
266  
267  
268  
269  
270  
271  
272  
273  
274  
275  
276  
277  
278  
279  
280  
281  
282  
283  
284  
285  
286  
287  
288  
289  
290  
291  
292  
293  
294  
295  
296  
297  
298  
299  
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  
233  
234  
235  
236  
237  
238  
239  
240  
241  
242  
243  
244  
245  
246  
247  
248  
249  
240  
241  
242  
243  
244  
245  
246  
247  
248  
249  
250  
251  
252  
253  
254  
255  
256  
257  
258  
259  
260  
261  
262  
263  
264  
265  
266  
267  
268  
269  
270  
271  
272  
273  
274  
275  
276  
277  
278  
279  
280  
281  
282  
283  
284  
285  
286  
287  
288  
289  
290  
291  
292  
293  
294  
295  
296  
297  
298  
299  
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  
233  
234  
235  
236  
237  
238  
239  
240  
241  
242  
243  
244  
245  
246  
247  
248  
249  
240  
241  
242  
243  
244  
245  
246  
247  
248  
249  
250  
251  
252  
253  
254  
255  
256  
257  
258  
259  
260  
261  
262  
263  
264  
265  
266  
267  
268  
269  
270  
271  
272  
273  
274  
275  
276  
277  
278  
279  
280  
281  
282  
283  
284  
285  
286  
287  
288  
289  
290  
291  
292  
293  
294  
295  
296  
297  
298  
299  
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  
233  
234  
235  
236  
237  
238  
239  
240  
241  
242  
243  
244  
245  
246  
247  
248  
249  
240  
241  
242  
243  
244  
245  
246  
247  
248  
249  
250  
251  
252  
253  
254  
255  
256  
257  
258  
259  
260  
261  
262  
263  
264  
265  
266  
267  
268  
269  
270  
271  
272  
273  
274  
275  
276  
277  
278  
279  
280  
281  
282  
283  
284  
285  
286  
287  
288  
289  
290  
291  
292  
293  
294  
295  
296  
297  
298  
299  
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  
233  
234  
235  
236  
237  
238  
239  
240  
241  
242  
243  
24

162 We see that parameters are updated with respect to features in two ways: for each feature  $j$ ,  $\nabla_{\theta_f} \alpha_j$   
 163 modifies the magnitude and frequency it fires, and  $\nabla_{\theta_f} e_j$  modifies the information carried. This  
 164 framework allows us to explain the following behaviors, through the lens of feature-based learning  
 165 dynamics.

166 **The model learns features with discrete, state-like presences rather than those with a continuous**  
 167 **spectrum of presences.** We see that the  $\alpha_j$  term controls the learning rate for  $\nabla_{\theta_f} e_j$ , suggesting  
 168 that models prefer to learn features  $e_j$  that naturally have high  $\alpha_j$ . (Section 4)

169 **High-impact features have phase transition points that form a centered, unimodal distribution.**  
 170 A transition point is where  $\frac{\partial \mathcal{L}}{\partial z}$  is highest, greatly affecting  $\nabla_{\theta_f} \alpha_j$ . The model is incentivized to  
 171 represent the presence of each feature in states that are far from this point. (Section 5.1)

172 **Low-impact features have phase transition points that form a symmetric bimodal distribution.**  
 173 Due to condensing high-dimensional information into low-dimensional space, feature presences are  
 174 affected by interference from superposition. While transition points would be centered between ideal  
 175 presences, the model minimizes damaging superposition by perturbing presences of low-impact  
 176 features. (Section 5.2)

### 179 3 APPLICATIONS TO 3D RECONSTRUCTION

181 To verify our analysis, we apply an SAE to Dora-VAE (Chen et al. 2024). Dora-VAE is a Variational  
 182 Autoencoder (VAE) that encodes point clouds  $\mathbf{P}_d$  sampled from 3D models to condensed latent  
 183 representations. These representations are then queried for diffusion-based reconstruction of the  
 184 initial geometry. Rather than a global latent for each shape, Dora-VAE selects a set  $\mathbf{P}_C$  of  $M$   
 185 point cloud features from  $\mathbf{P}_d$  using furthest point sampling (FPS), which is passed through several  
 186 cross-attention layers alongside  $\mathbf{P}_d$ . This forms a set  $\mathbf{C}$  of processed point features.

$$187 \mathbf{P}_C = \text{FPS}(\mathbf{P}_d) \quad (7)$$

$$188 \mathbf{C} = \text{CrossAttn}(\text{PosEnc}(\mathbf{P}_C), \text{PosEnc}(\mathbf{P}_d))$$

190 We take the provided Dora-VAE network, pretrained on a subset of Objaverse, and encode 53k  
 191 objects from Objaverse-XL. These encodings form a dataset of pre-KL embedding network states.  
 192 The number of latents in  $\mathbf{C}$  is determined by the number of points initially sampled; we record a set  
 193 for  $M = 4096$ , where each pre-embedding is size 128.

194 After encoding, for each pre-embedding, Dora-VAE isolates a mean  $\mu_i \in \mathbb{R}^{64}$  and variance  $\sigma_i \in$   
 195  $\mathbb{R}^{64}$  by chunking.

$$196 \mathbf{C} = \{(\mu_i, \sigma_i)\}_{i=1}^M$$

197 Thus,  $\forall i \in \{1, 2, \dots, M\}, j \in \{1, 2, \dots, 64\}$ , the KL embedding is:

$$198 z_{i,j} = \mu_{i,j} + \sigma_{i,j} \cdot \epsilon \quad (8)$$

200 where  $\epsilon \sim \mathcal{N}(0, 1)$ . This embedding is fed through the decoder before querying for occupancy.  
 201 Here, the term *latent* with respect to Dora-VAE will refer to the network state post-KL embedding.

#### 204 3.1 SAE ON DORA-VAE

206 The dataset for our SAE is constructed from these recorded pre-embeddings. Each epoch, through  
 207 KL, we sample each recorded pre-embedding for new latents. Our latent space is thus extremely  
 208 well-defined, as each epoch of training has 217 million newly sampled latents. In addition, since  
 209 these latents are point cloud features initially downsampled from  $\mathbf{P}_d$ , each latent will correspond to  
 210 a point of the initial point cloud sample. This relationship allows us to interpret a feature based off  
 211 of the position or structure of points with high presence for that feature.

212 We train our BatchTopK SAE with  $M = 4096$ , codebook size  $n = 512$ , threshold  $k = 8$ , and  
 213  $\beta = 0.125$ . We use a batch size of 327680 latents, randomly selected regardless of which 3D model  
 214 produced each latent. We use the Adam optimizer with an initial learning rate of 1e-3 and train for  
 215 ten epochs. The model was trained on a single A100 and took 2 hours to train. We also present  
 metrics for variations on codebook size and threshold in Appendix B.1.

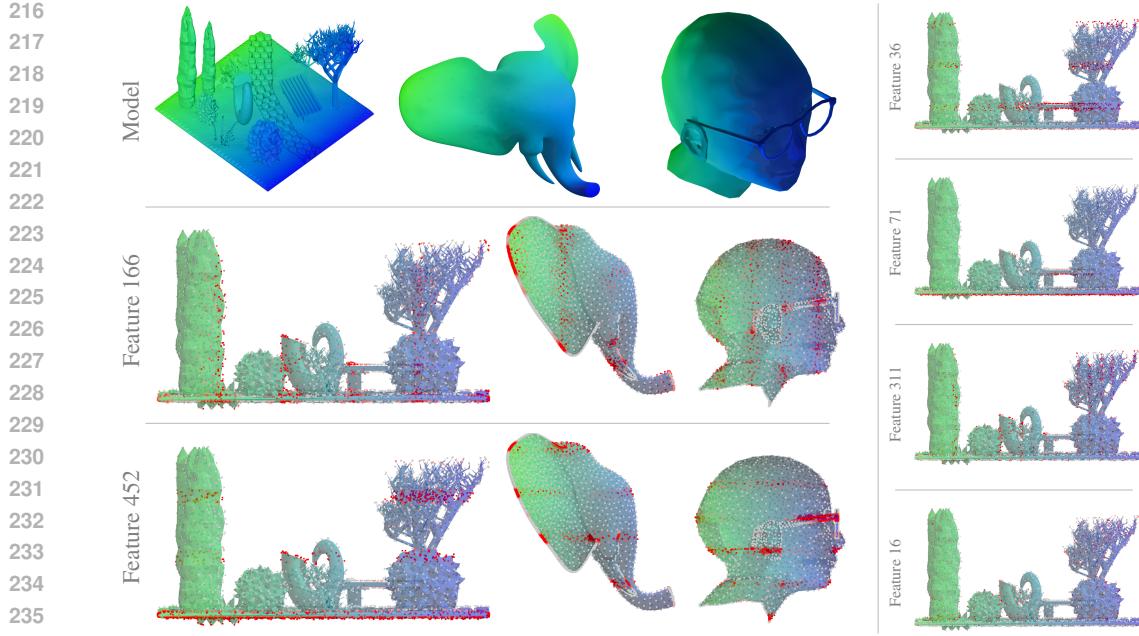


Figure 2: Overview of Dora-VAE features. Each point’s color represents the presence of feature  $j$  at that point. Visually, most features congregate in stripes along a single axis. This suggests that continuous position is represented by a set of discrete elements, which each activate in separate regions. We show further examples in Appendix C.

We highlight the qualitative performance of our feature extraction in Figure 2. For each encoded 3D object, we obtain a set of latents  $\{\mathbf{z}_i\}_{i=1}^M$ , and  $\alpha_i$  for each latent by passing it through the SAE. We plot the  $M$  latents as points from their initially sampled positions  $\mathbf{P}_d$ . Finally, to examine a feature  $j$ , we color each point  $i$  of each latent based off the presence  $\alpha_{i,j}$ .

Most features display positional information along a single axis. Notably, features appear state-like, and store information in a binary manner. Features emerge at striped intervals across models in a manner akin to positional encoding, suggesting latents form a discrete representation. In other words, a feature doesn’t have a range of possible values (“As feature  $j$  increases, the point travels further along the axis”) — instead, a feature tends towards one of two states (“If feature  $j$  is present, it is within this region”). We discuss the discretization of features further in Section 4.

These positional features are highly visually interpretable due to the 3D medium. We see that such features are applicable across all models, and activate with significant sensitivity and specificity. Some features, although they are highly present across the model, have meanings that are difficult to interpret through observation. We can instead intervene on these features to determine their purpose.

### 3.2 FEATURE ABLATION AND ADDITION

It is possible these features are simply vestiges of correlations between the sampled points; points that share close coordinates may simply propagate similarly across the encoder. To disprove this, and demonstrate these features are meaningful internal representations, we examine the downstream effects of modifying latents along feature axes.

We intervene on features through ablation and addition based on SAE decoder weights. In our pipeline, inputs are encoded by Dora-VAE to a set of latent vectors  $\{\mathbf{z}_i\}_{i=1}^M$ . We recall Eq. 2; to visualize the effect of modifying feature  $j$  on the reconstruction, we want to approximate a modified set of latents such that:

$$\mathbf{z}'_i = \mathbf{E}^T \boldsymbol{\alpha}'_i \quad \text{where} \quad \boldsymbol{\alpha}'_i = \boldsymbol{\alpha}_i, \alpha'_{i,j} = (1 - t) \cdot \alpha_{i,j} \quad (9)$$

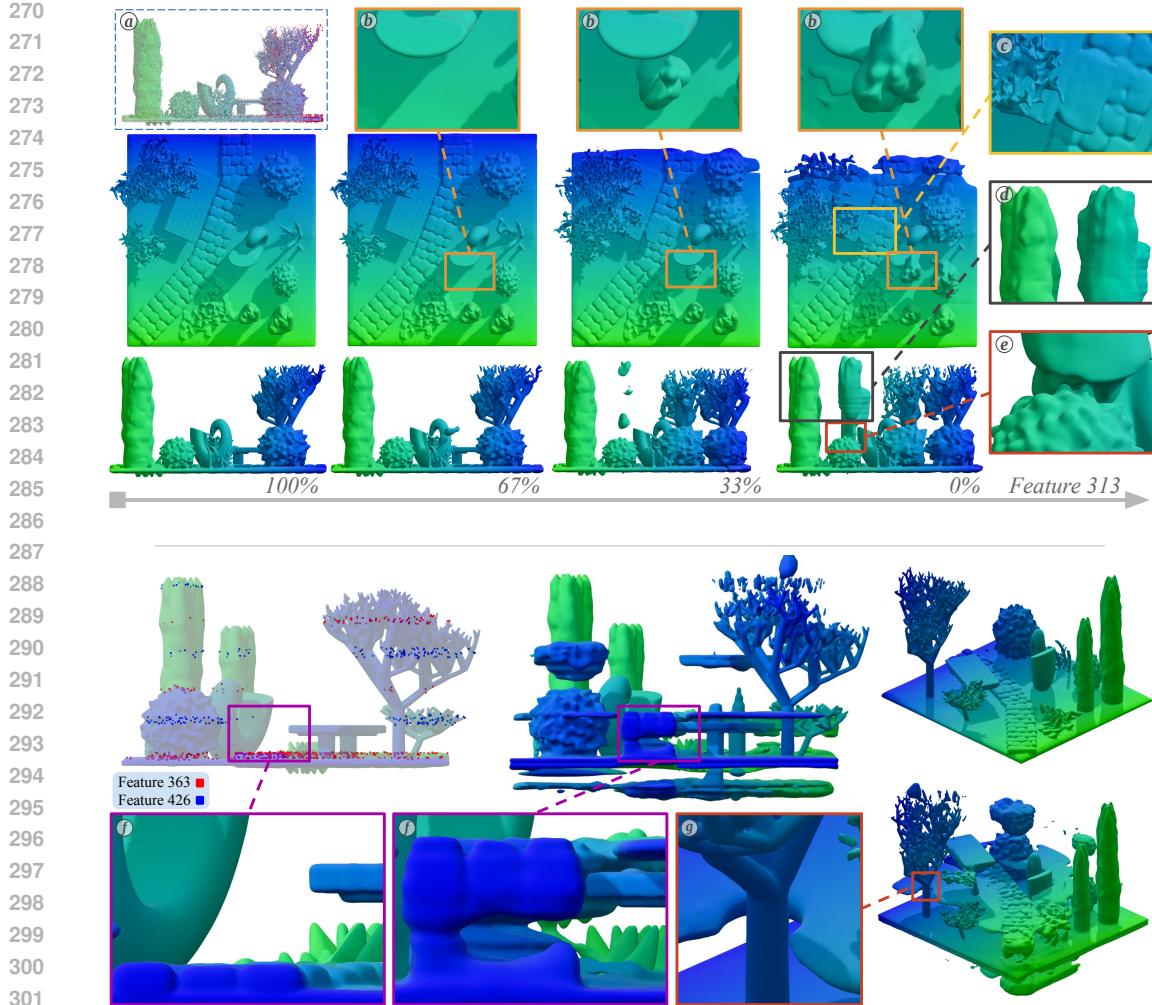


Figure 3: *Top*: Example ablation on feature 313. (a) This feature primarily attends to points on the positive end of the  $z$ -axis. As the feature is removed, shapes disappear and appear spontaneously, rather than moving along the object (b). This suggests the feature represents a discrete region in space, rather than a continuous range of positions. *Bottom*: Example of feature 363 ablation and feature 426 addition. Both features attend to regions along the  $y$ -axis. (f) Shapes that had their region fixed by feature 363 are moved to regions defined by feature 426. In addition, they preserve their local structure. See Appendix B.2.

where  $(1 - t)$  is the proportion of the original presence, set externally. Rather than rely on the reconstruction provided by our SAE, we modify the latents with the decoder weight for feature  $j$ .

$$\forall \mathbf{z}_i \in \{\mathbf{z}_i\}_{i=1}^M \left\{ \begin{array}{ll} \text{Ablation: } \mathbf{z}'_i \approx \mathbf{z}_i - t \cdot \text{Enc}(\mathbf{z}_i)_j \mathbf{w}_j^{\text{dec}} \\ \text{Addition: } \mathbf{z}'_i \approx \mathbf{z}_i + \alpha'_j \mathbf{w}_j^{\text{dec}} \end{array} \right. \quad (10)$$

During ablation, the portion of feature  $j$  removed is scaled by value  $t$ ; here,  $t = 1$  implies the feature is completely removed. During addition, feature  $j$  is added by amount of a manually set  $\alpha'_j$ . Each modified set of latents is passed through the the Dora-VAE decoder to be compared to the original model. We record the mean squared error (MSE) of the decoded reconstruction.

The top of Figure 3 demonstrates an ablation of feature 313. As shown by the red points in (a), this feature primarily attends to points on the positive end of the  $z$ -axis, with a small region on the negative end. When the feature is ablated, shapes whose position relied on it are rendered elsewhere — this is indicative of a causal relationship between this feature and the shape's position. In addition,

324 as shown in ⑥, rendered points appear spontaneously, rather than moving across the model. This  
 325 again suggests that features represent discrete states, and presences do not have a continuous range  
 326 of information.

327 At the bottom of Figure 3, we demonstrate an ablation of feature 363 alongside an addition of feature  
 328 426. Rather than applying a constant presence of feature 426 on all latents, we instead replace every  
 329 presence of feature 363 with an equal presence of feature 426.

$$331 \quad \forall \mathbf{z}_i \in \{\mathbf{z}_i\}_{i=1}^M, \quad \mathbf{z}'_i \approx \mathbf{z}_i - \text{Enc}(\mathbf{z}_i)_{363} \mathbf{w}_{363}^{dec} + \text{Enc}(\mathbf{z}_i)_{363} \mathbf{w}_{426}^{dec}$$

333 Note that, because both features 363 and 426 attend to positions on the y-axis, points displaced by  
 334 the removal of feature 363 are anchored by the addition of feature 426. Rendered shapes preserve  
 335 their form even after moving. We discuss further observations in Appendix B.2.

## 337 4 DO FEATURES SHOW STATE-BASED BEHAVIOR?

339 A significant portion of features in Dora-VAE are dedicated to representing the position of the la-  
 340 tent. Intuitively, one would assume that these features should be continuous, as points are relatively  
 341 uniformly distributed across 3D space. Each latent’s position could be represented by only three  
 342 features, with others for additional fidelity. Despite this, the model chooses to represent features  
 343 discretely; if the feature has a high presence, the position is within a defined region. This method of  
 344 representation is akin to binary positional encoding. Some features make the similarity more explicit  
 345 by representing a set of multiple regions across an axis, rather than a single one.

346 To verify whether these features are truly discrete, we perform a series of systematic feature ab-  
 347 lations over our dataset and measure the change in loss. As above, we pass the  $M$  latents of 53k  
 348 3D objects through our BatchTopK SAE with  $k = 8$  and codebook size 512. **Each 3D object thus**  
 349 **has a set of presence vectors  $\{\alpha_i\}_{i=1}^M$  where each  $\alpha_i \in \mathbb{R}^{512}$  shows the presences of 512 fea-**  
 350 **tures and has 8 nonzero values on average.** For each 3D object, we randomly select 16 features to  
 351 intervene on, preferring features that are present in more latents. We perform each ablation with  
 352  $t \in \{0.00, 0.05, 0.10, \dots, 1.0\}$ , recording the MSE of the decoded reconstruction for each  $t$ . We  
 353 evaluated 848k ablations in total.

354 The model’s response to feature ablation displays interesting recurring behaviors. **Given our set**  
 355  **$\{\alpha_i\}_{i=1}^M$ , we define the feature density of feature  $j$  as  $\frac{1}{M} \sum_{i=1}^M \mathbf{1}\{\alpha_{i,j} \neq 0\}$ , and the average pres-**  
 356 **ence of feature  $j$  as  $\frac{1}{M} \sum_{i=1}^M \alpha_{i,j}$ . We also define the impact  $\Delta L$  of an ablation as the difference**  
 357 **in MSE between  $t = 0$  and  $t = 1$ .** Figure 4 shows kernel density estimations (KDEs) for these  
 358 properties of each feature ablation. We note that our set of ablations shows a wide variety of im-  
 359 **port, feature density, and average presence.** In addition, impact is positively correlated with both the  
 360 **feature density and average presence.**

361 We also plot several ablation-response curves (ARCs). **Each curve represents a single ablation, and**  
 362 **shows the change in MSE as  $t$  increases. We normalize MSE such that the plotted error at  $t = 0$**   
 363 **is 0, and at  $t = 1$  is 1. We also record the transition point of an ablation as the value of  $t$  when**  
 364 **the normalized MSE is 0.5.** Note that the ARCs do not show a linear relationship between change  
 365 in latents and MSE. Rather, they exhibit variable curvature, with two inflection points — initial  
 366 changes in loss are below our projected linear growth, then accelerate at the transition point, before  
 367 again slowing down.

368 We find ARCs with greater impact exhibit more discrete behavior. To demonstrate this, we group  
 369 ARCs together based on  $\Delta L$  and perform four experiments. **First**, for each group of similar  $\Delta L$ , we  
 370 plot a KDE of all normalized MSE for  $0.05 \leq t \leq 0.95$ . Notably, as impact increases, intermediate  
 371 MSE values tend to cluster towards the initial and final MSE. **Second**, we analyze the maximum  
 372 slope of each ARC to determine if it is an outlier in the distribution of slopes. For each ARC, we  
 373 estimate slope as the difference of normalized MSE between every two consecutive  $t$  values. We  
 374 then z-score the greatest slope of each ARC relative to the distribution formed by all slopes in the  
 375 group, and plot the average for each group. As the  $\Delta L$  of an ARC increases, the greatest slope of the  
 376 ARC trends further from the group distribution, suggesting the transition point is more well defined  
 377 as impact increases. **Third**, we find the point of flattest slope for each ablation, record the value of  
 $t$ , and plot a KDE of these values. We see that points near the beginning and end have typically the

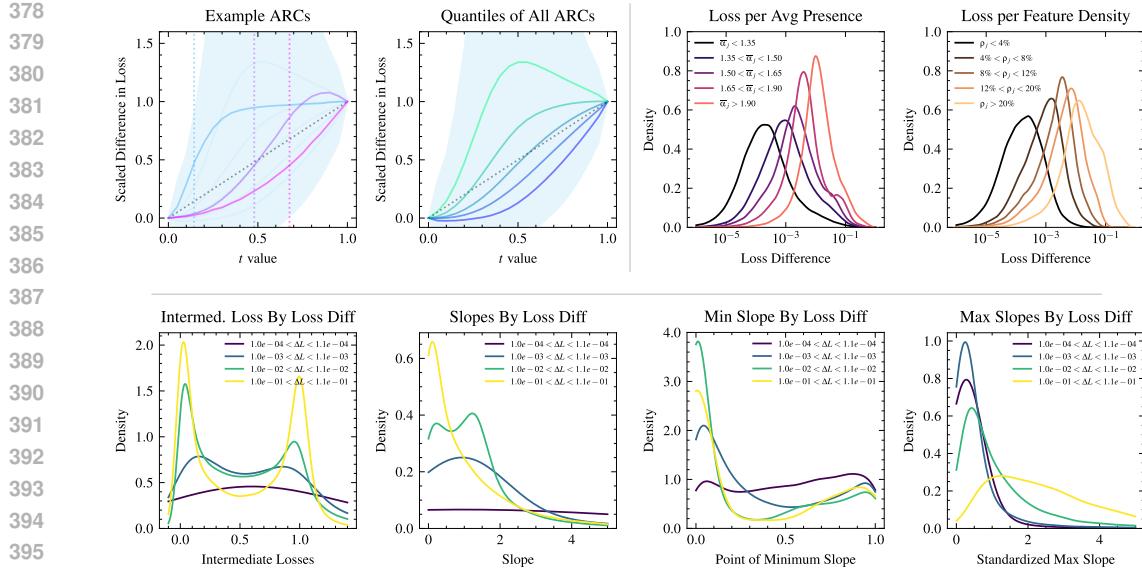


Figure 4: *Top left*: Examples and quantiles of ARCs. ARCs display an almost sigmoidal behavior, with a varied transition point. *Top right*: Correlations of loss with average value and feature density. *Bottom*: Various experiments to demonstrate discretization. ARCs will typically have stagnant MSE near the beginning and end, and change most rapidly in a small interval.

flattest slope. Finally, we note that the distribution of slopes across ARCs leans further left as  $\Delta L$  increases, showing that ARCs are typically flatter, with sharper jumps, when impact increases.

The discretization of features can be explained through the learning dynamics defined in Section 2.1. The signal to the identity of feature  $j$ ,  $\nabla_{\theta_f} e_j$ , is scaled by the presence  $\alpha_j$ . Thus, the identity of feature  $j$  is most influenced when  $\alpha_j$  is high, while, when  $\alpha_j$  is low, the signal to identity is diluted by other signals for features with higher presence. We further discuss this intuition in Appendix B.3.

## 5 THE BIMODALITY OF TRANSITION POINTS

If we interpret feature activations as a discrete state space with distinct phase transitions, we can follow up by investigating when these phase transitions occur. We plot KDEs of two properties — transition points and the points of greatest slope — for all ARCs in Figure 5. Surprisingly, both distributions are bimodal.

We further investigate this behavior by again grouping ARCs by  $\Delta L$  and plotting the KDEs of transition points and points of greatest slope for each group. From this figure, we see that the transition points of high-impact ARCs form a unimodal distribution around the center ( $t \approx 0.5$ ), while the transition points of low-impact ARCs form a bimodal distribution roughly symmetric about this center.

We then further group ablations based on which feature  $j$  is removed and repeat the same investigations as above. Again, for high-impact ablations, each feature has a peak near the center. However, low-impact ablations are no longer bimodal, and instead there is instead a single peak that strays from the center. Some fall closer to the beginning, while others are nearer the end. Thus, it is only when all ARCs are aggregated together, regardless of feature, that we observe a bimodality of transition points in low-impact features.

### 5.1 UNIMODAL TRANSITION POINTS OF HIGH-IMPACT ABLATIONS

We explain the distribution of high-impact transition points using the learning dynamics defined in Equation 6. Because our feature activations, especially high-impact ones, approximate discrete

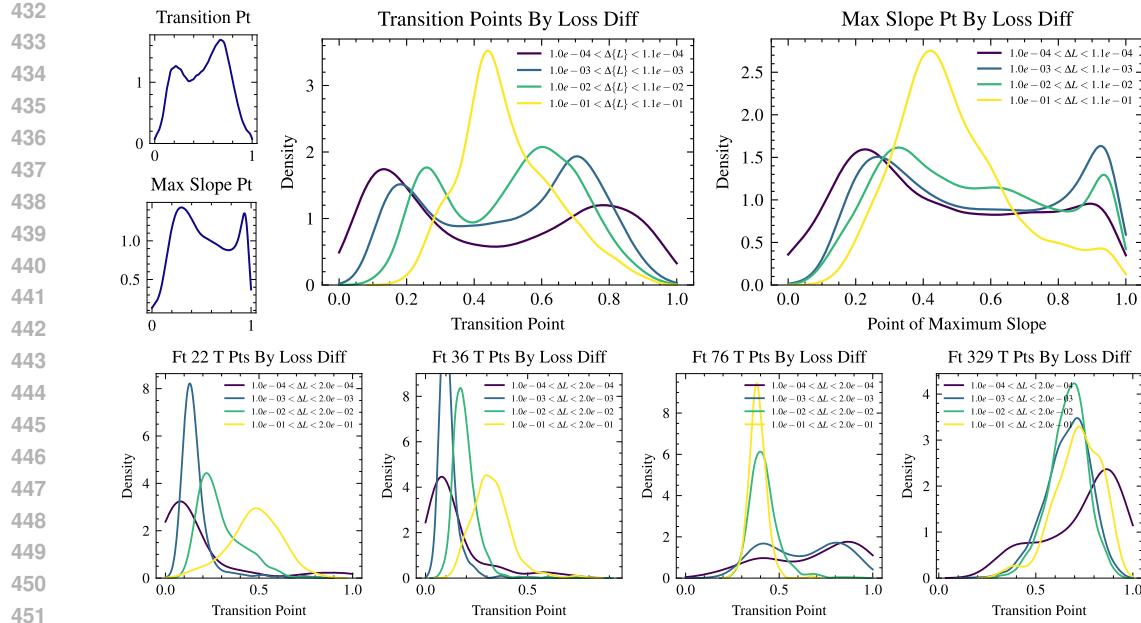


Figure 5: *Top*: Distributions of transition points and points of maximum slope over all ARCs, grouped by  $\Delta L$ . ARCs with greater  $\Delta L$  form a unimodal distribution of transition points centered about  $t \approx 0.5$ , while those with less  $\Delta L$  have a symmetric, bimodal distribution. *Bottom*: Distributions of transition points of individual features, grouped by  $\Delta L$ . ARCs with greater  $\Delta L$  still form a somewhat centered, unimodal distribution, but those with less  $\Delta L$  now also unimodal, further from the center. This polarization across individual features causes the bimodal distribution when considered in aggregate.

behavior, we can consider a feature to be *on* (high presence) or *off* (low presence). During the gradient step, the effect on  $\nabla_{\theta_f} \alpha_j$  is scaled by  $\frac{\partial \mathcal{L}}{\partial z}$ . This gradient, by definition, reaches a peak near the transition point. Because of the rapid change in loss associated with the transition point,  $\nabla_{\theta_f} \alpha_j$  is incentivized to adjust  $\alpha_j$  such that both on and off states of the feature are at a distance from the transition point. The magnitude of  $\frac{\partial \mathcal{L}}{\partial z}$  at the transition point is not necessarily equivalent for when feature  $j$  is on and off — however, over many different transition points, we can assume it is roughly symmetric. Thus, the overall distribution of high-impact transition points is located at the center.

## 5.2 BIMODAL TRANSITION POINTS OF LOW-IMPACT ABLATIONS

The bimodality of low-impact transition points is a more complex property. When examining individual features, the distribution of low-impact transition points forms a unimodal peak that drifts away from the center. This behavior is not caused by a weaker  $\frac{\partial L}{\partial z}$  at the transition point, which would've only increased the variance of the distribution. Instead, the peak itself is offset to the left or right — as if a polarizing effect drives transition points away from the center as  $\Delta L$  decreases.

We speculate that this effect is caused by a variable offset applied to the feature presence, and provide a visual aid in Figure 6. We suggest that as impact decreases, the larger this offset becomes. In this way, the transition point is moved earlier or later in the ablation. This effect does not make the feature presence estimation less accurate; again, in that case, we would see the low-impact distribution have greater variance, but retain the same center. Rather, the peak itself moves. We hypothesize that this offset is caused by the model learning to redistribute interference from superposition.

Superposition occurs when high-dimensional features are constrained to a low dimensional space, causing interference between features in previously distinct dimensions. (Elhage et al. 2022). We show another visual aid in Figure 7. Suppose the model has determined to represent a latent through high-impact feature 1, low-impact feature 2, and a set of several other features  $*$ . We refer to the presences as  $\alpha_1$ ,  $\alpha_2$ , and  $\alpha_*$ . As shown at the top of Figure 7, due to superposition, the features  $*$

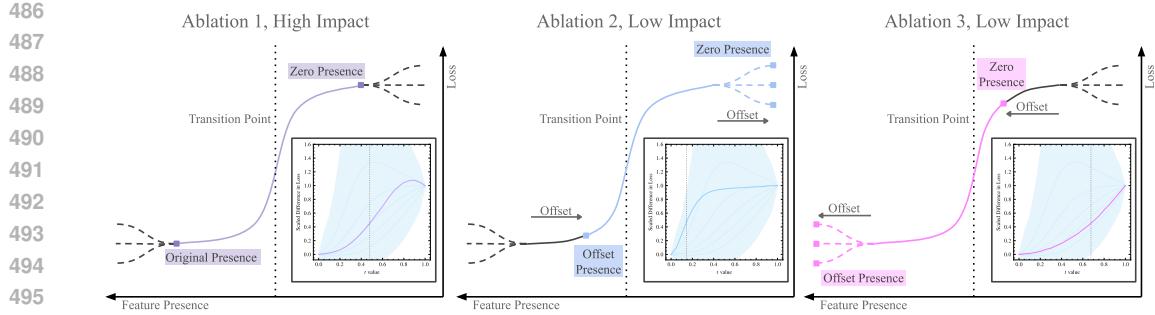


Figure 6: *Left*: A visual aid of a high-impact ARC. As feature presence is ablated, the loss (MSE) increases. It is fastest when passing through the transition point, then stagnates. *Center and Right*: A visual aid of two low-impact ARCs. We suggest that low-impact ARCs are affected by a polarizing offset. If the presence is decreased by the offset, the relative transition point is moved left. If increased, the point is moved right.

add interference in the direction of  $\alpha_1$  and  $\alpha_2$ . We suggest that the model selects features  $*$  such that feature 2, a low-impact feature, is affected by this superposition more than feature 1, a high-impact feature. This causes the offset shown at the bottom of Figure 7 — while the transition point of feature 1 is still centered, the transition point of feature 2 is moved.

Even though  $\Delta L$  has correlations with other variables, they are not strong enough to cause a third variable contamination that affects SAE estimations of presences (see Appendix B.4). Thus, we can reasonably conclude the offset is present even before the latent is passed to the SAE. In other words, we speculate the model itself is actively preserving high-impact features by passing interference to features that are lower impact.

Previous works have shown that a model learns a set of features that minimizes interference from superposition during training. However, this work suggests that a model can redistribute interference from superposition *at inference time*; that not only does the model select the most important features with respect to the dataset, but it dynamically estimates the importance of features *input to input*.

## 6 FUTURE WORK AND CONCLUSION

These results are very motivating for model learning dynamics research, but further work is necessary to generalize and substantiate our findings. First, by validating our observations across other domains (text, image) and models (PointNet++ (Qi et al. 2017), LION (Zeng et al. 2022)), we can confirm our results here and broaden the scope of our work. Second, examining Dora-VAE with circuit detection techniques (Ameisen et al. 2025) and extracting an attribution graph may show how features flow through the architecture and form discrete patterns. Third, probing gradients of toy models, similar to Elhage et al. (2022), will evaluate the feature learning framework we present here. Finally, if we can identify what influences redistribution of interference among features, we can potentially perform feature decomposition at training time to develop a meta-learning module.

Our work is the first to apply an SAE to 3D data, highlighting specific discovered features and showing the causally related downstream effects. We then take advantage of the continuous and unstructured nature of the domain to investigate the model’s feature decomposition, confirming that the latent space can be interpreted as a discrete, state-based feature space driven by phase transitions. We then provide a potentially general framework of feature learning dynamics that explains the unexpected discretization we observe. Finally, we explain the counter-intuitive property of the bimodality of transition points by proposing a mechanism by which the model redistributes superposition to only affect low-impact features.

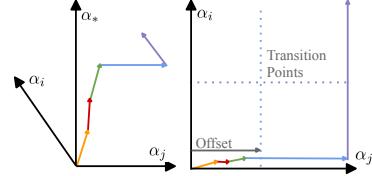


Figure 7: Superposition interference affects feature presences. We suggest this offsets low-impact features, shifting the relative transition point.

540 ACKNOWLEDGMENTS  
541542 Acknowledgments have been anonymized for review.  
543544 REFERENCES  
545

546 Emmanuel Ameisen, Jack Lindsey, Adam Pearce, Wes Gurnee, Nicholas L. Turner, Brian Chen,  
547 Craig Citro, David Abrahams, Shan Carter, Basil Hosmer, Jonathan Marcus, Michael Sklar,  
548 Adly Templeton, Trenton Bricken, Callum McDougall, Hoagy Cunningham, Thomas Henighan,  
549 Adam Jermyn, Andy Jones, Andrew Persic, Zhenyi Qi, T. Ben Thompson, Sam Zimmerman,  
550 Kelley Rivoire, Thomas Conerly, Chris Olah, and Joshua Batson. Circuit tracing: Revealing  
551 computational graphs in language models. *Transformer Circuits Thread*, 2025. URL <https://transformer-circuits.pub/2025/attribution-graphs/methods.html>.  
552

553 Yoshua Bengio, Aaron Courville, and Pascal Vincent. Representation learning: A review and new  
554 perspectives. *IEEE Transactions on Pattern Analysis and Machine Intelligence*, 35(8):1798–1828,  
555 2013. doi: 10.1109/TPAMI.2013.50. URL <https://doi.org/10.1109/TPAMI.2013.50>.  
556

557 Trenton Bricken, Adly Templeton, Joshua Batson, Brian Chen, Adam Jermyn, Tom Conerly,  
558 Nick Turner, Cem Anil, Carson Denison, Amanda Askell, Robert Lasenby, Yifan Wu,  
559 Shauna Kravec, Nicholas Schiefer, Tim Maxwell, Nicholas Joseph, Zac Hatfield-Dodds, Alex  
560 Tamkin, Karina Nguyen, Brayden McLean, Josiah E Burke, Tristan Hume, Shan Carter,  
561 Tom Henighan, and Christopher Olah. Towards monosemanticity: Decomposing language  
562 models with dictionary learning. *Transformer Circuits Thread*, 2023. <https://transformer-circuits.pub/2023/monosemantic-features/index.html>.  
563

564 Bart Bussmann, Patrick Leask, and Neel Nanda. Batchtopk sparse autoencoders, 2024. URL  
565 <https://arxiv.org/abs/2412.06410>.  
566

567 Rui Chen, Jianfeng Zhang, Yixun Liang, Guan Luo, Weiyu Li, Jiarui Liu, Xiu Li, Xiaoxiao Long,  
568 Jia Shi Feng, and Ping Tan. Dora: Sampling and benchmarking for 3d shape variational auto-  
569 encoders. *arXiv preprint arXiv:2412.17808*, 2024. doi: 10.48550/arXiv.2412.17808. URL  
570 <https://arxiv.org/abs/2412.17808>. Accepted to CVPR 2025.  
571

572 Hoagy Cunningham, Aidan Ewart, Logan Riggs, Robert Huben, and Lee Sharkey. Sparse autoen-  
573 coders find highly interpretable features in language models. In *Proceedings of the International  
574 Conference on Learning Representations (ICLR 2024)*, 2024. doi: 10.48550/arXiv.2309.08600.  
575 URL <https://arxiv.org/abs/2309.08600>.  
576

577 Matt Deitke, Dustin Schwenk, Jordi Salvador, Luca Weihs, Oscar Michel, Eli VanderBilt, Ludwig  
578 Schmidt, Kiana Ehsani, Aniruddha Kembhavi, and Ali Farhadi. Objaverse: A universe of anno-  
579 tated 3d objects. *arXiv preprint arXiv:2212.08051*, 2022. doi: 10.48550/arXiv.2212.08051. URL  
580 <https://arxiv.org/abs/2212.08051>.  
581

582 Nelson Elhage, Tristan Hume, Catherine Olsson, Nicholas Schiefer, Tom Henighan, Shauna Kravec,  
583 Zac Hatfield-Dodds, Robert Lasenby, Dawn Drain, Carol Chen, Roger Grosse, Sam McCandlish,  
584 Jared Kaplan, Dario Amodei, Martin Wattenberg, and Christopher Olah. Toy models of super-  
585 position. *Transformer Circuits Thread*, 2022. URL [https://transformer-circuits.pub/2022/toy\\_model/index.html](https://transformer-circuits.pub/2022/toy_model/index.html).  
586

587 Nelson Elhage, Tristan Hume, Catherine Olsson, Nicholas Schiefer, Tom Henighan, ..., and  
588 Christopher Olah. Superposition, memorization, and double descent. *Transformer Circuits  
589 Thread*, 2023. <https://transformer-circuits.pub/2023/toy-double-descent/index.html>.  
590

591 Thomas Fel, Victor Boutin, Mazda Moayeri, Rémi Cadène, Louis Bethune, Léo Andéol, Math-  
592 ieu Chalvidal, and Thomas Serre. A holistic approach to unifying automatic concept ex-  
593 traction and concept importance estimation. *arXiv preprint arXiv:2306.07304*, 2023a. URL  
594 <https://arxiv.org/abs/2306.07304>.

594 Thomas Fel, Agustin Picard, Louis Bethune, Thibaut Boissin, David Vigouroux, Julien Colin, Rémi  
 595 Cadène, and Thomas Serre. Craft: Concept recursive activation factorization for explainability. In  
 596 *Proceedings of the IEEE/CVF Conference on Computer Vision and Pattern Recognition (CVPR)*,  
 597 pp. 18526–18536, 2023b.

598 Thomas Fel, Ekdeep Singh Lubana, Jacob S. Prince, Matthew Kowal, Victor Boutin, Isabel Pa-  
 599 padimitriou, Binxu Wang, Martin Wattenberg, Demba Ba, and Talia Konkle. Archetypal sae:  
 600 Adaptive and stable dictionary learning for concept extraction in large vision models, 2025. URL  
 601 <https://arxiv.org/abs/2502.12892>.

602 Leo Gao, Tom DuprélaTour, Henk Tillman, Gabriel Goh, Rajan Troll, Alec Radford, Ilya  
 603 Sutskever, Jan Leike, and Jeffrey Wu. Scaling and evaluating sparse autoencoders. *Open-  
 604 AI Technical Report*, 2025. Preprint available at <https://cdn.openai.com/papers/sparse-autoencoders.pdf>.

605 Edith Natalia Villegas Garcia and Alessio Ansuini. Interpreting and steering protein language mod-  
 606 els through sparse autoencoders, 2025. URL <https://arxiv.org/abs/2502.09135>.

607 Amirata Ghorbani, James Wexler, James Y Zou, and Been Kim. Towards automatic concept-based  
 608 explanations. In *Advances in Neural Information Processing Systems*, volume 32, pp. 9273–  
 609 9282, 2019. URL [https://papers.nips.cc/paper\\_files/paper/2019/hash/77d2afcb31f6493e350fca61764efb9a-Abstract.html](https://papers.nips.cc/paper_files/paper/2019/hash/77d2afcb31f6493e350fca61764efb9a-Abstract.html).

610 Liv Gorton and Owen Lewis. Adversarial examples are not bugs, they are superposition, 2025. URL  
 611 <https://arxiv.org/abs/2508.17456>.

612 Kaarel Hänni, Jake Mendel, Dmitry Vaintrob, and Lawrence Chan. Mathematical models of com-  
 613 putation in superposition, 2024. URL <https://arxiv.org/abs/2408.05451>.

614 Been Kim, Martin Wattenberg, Justin Gilmer, Carrie Cai, James Wexler, Fernanda Viégas, and  
 615 Rory Sayres. Interpretability beyond feature attribution: Quantitative testing with concept acti-  
 616 vation vectors (tcav). In Jennifer Dy and Andreas Krause (eds.), *Proceedings of the 35th Inter-  
 617 national Conference on Machine Learning*, volume 80 of *Proceedings of Machine Learning Re-  
 618 search*, pp. 2668–2677. PMLR, 2018. URL <https://proceedings.mlr.press/v80/kim18d.html>.

619 Jack Lindsey, Wes Gurnee, Emmanuel Ameisen, Brian Chen, Adam Pearce, Nicholas L. Turner,  
 620 Craig Citro, David Abrahams, Shan Carter, Basil Hosmer, Jonathan Marcus, Michael Sklar, Adly  
 621 Templeton, Trenton Bricken, Callum McDougall, Hoagy Cunningham, Thomas Henighan, Adam  
 622 Jermyn, Andy Jones, Andrew Persic, Zhenyi Qi, T. Ben Thompson, Sam Zimmerman, Kelley  
 623 Rivoire, Thomas Conerly, Chris Olah, and Joshua Batson. On the biology of a large language  
 624 model. *Transformer Circuits Thread*, 2025. URL <https://transformer-circuits.pub/2025/attribution-graphs/biology.html>.

625 Francesco Locatello, Stefan Bauer, Mario Lucic, Gunnar Rätsch, Sylvain Gelly, Bernhard  
 626 Schölkopf, and Olivier Bachem. Challenging common assumptions in the unsupervised learn-  
 627 ing of disentangled representations. In *Proceedings of the 36th International Conference on Ma-  
 628 chine Learning*, pp. 4114–4124, 2019. URL <https://proceedings.mlr.press/v97/locatello19a.html>.

629 Samuel Marks, Tegan Chevalier, Jacob Steinhardt, Chris Olah, Rishi Bommasani, et al. Sparse  
 630 feature circuits: Discovering and editing interpretable causal graphs in language models. *arXiv  
 631 preprint arXiv:2403.19647*, 2024. URL <https://arxiv.org/abs/2403.19647>.

632 Tomas Mikolov, Wen-tau Yih, and Geoffrey Zweig. Linguistic regularities in continuous space  
 633 word representations. In *Proceedings of the 2013 Conference of the North American Chapter  
 634 of the Association for Computational Linguistics: Human Language Technologies*, pp. 746–751,  
 635 2013. URL <https://aclanthology.org/N13-1090>.

636 Core Francisco Park, Maya Okawa, Andrew Lee, Ekdeep Singh Lubana, and Hidenori Tanaka.  
 637 Emergence of hidden capabilities: Exploring learning dynamics in concept space. In *Ad-  
 638 vances in Neural Information Processing Systems 37 (NeurIPS 2024)*, June 2024. doi: 10.  
 639 48550/arXiv.2406.19370. URL <https://arxiv.org/abs/2406.19370>. arXiv preprint  
 640 arXiv:2406.19370.

648 Charles R. Qi, Li Yi, Hao Su, and Leonidas J. Guibas. Pointnet++: Deep hierarchical feature  
 649 learning on point sets in a metric space, 2017. URL <https://arxiv.org/abs/1706.02413>.

650

651 Senthooran Rajamanoharan, Tom Lieberum, Nicolas Sonnerat, Arthur Conmy, Vikrant Varma, János  
 652 Kramár, and Neel Nanda. Jumping ahead: Improving reconstruction fidelity with jumprelu sparse  
 653 autoencoders, 2024. URL <https://arxiv.org/abs/2407.14435>.

654

655 Sukrut Rao, Sweta Mahajan, Moritz Böhle, and Bernt Schiele. Discover-then-name: Task-agnostic  
 656 concept bottlenecks via automated concept discovery, 2024. URL <https://arxiv.org/abs/2407.14499>.

657

658 Adly Templeton, Tom Conerly, Jonathan Marcus, Jack Lindsey, Trenton Bricken, Brian Chen,  
 659 Adam Pearce, Craig Citro, Emmanuel Ameisen, Andy Jones, Hoagy Cunningham, Nicholas L  
 660 Turner, Callum McDougall, Monte MacDiarmid, C. Daniel Freeman, Theodore R. Sumers,  
 661 Edward Rees, Joshua Batson, Adam Jermyn, Shan Carter, Chris Olah, and Tom Henighan.  
 662 Scaling monosemanticity: Extracting interpretable features from claude 3 sonnet. *Transformer*  
 663 *Circuits Thread*, 2024. URL <https://transformer-circuits.pub/2024/scaling-monosemanticity/index.html>.

664

665 Harrish Thasarathan, Julian Forsyth, Thomas Fel, Matthew Kowal, and Konstantinos Derpanis. Uni-  
 666 versal sparse autoencoders: Interpretable cross-model concept alignment, 2025. URL <https://arxiv.org/abs/2502.03714>.

667

668 Eric Todd, Millicent L. Li, Arnab Sen Sharma, Aaron Mueller, Byron C. Wallace, and David Bau.  
 669 Function vectors in large language models. In *Proceedings of the 2024 International Conference*  
 670 *on Learning Representations*, 2024. URL <https://openreview.net/forum?id=AwyxtyMwaG>. arXiv:2310.15213.

671

672 Hanqi Yan, Xiangxiang Cui, Lu Yin, Paul Pu Liang, Yulan He, and Yifei Wang. The multi-faceted  
 673 monosemanticity in multimodal representations. *arXiv preprint arXiv:2502.14888*, 2025. URL  
 674 <https://arxiv.org/abs/2502.14888>.

675

676 Xiaohui Zeng, Arash Vahdat, Francis Williams, Zan Gojcic, Or Litany, Sanja Fidler, and Karsten  
 677 Kreis. Lion: Latent point diffusion models for 3d shape generation, 2022. URL <https://arxiv.org/abs/2210.06978>.

678

679 Ruihan Zhang, Prashan Madumal, Tim Miller, Krista A. Ehinger, and Benjamin I. P. Rubinstein.  
 680 Invertible concept-based explanations for cnn models with non-negative concept activation vec-  
 681 tors. In *Proceedings of the Thirty-Fifth AAAI Conference on Artificial Intelligence (AAAI 2021)*,  
 682 pp. 11682–11690, 2021. doi: 10.1609/aaai.v35i13.17389. URL <https://arxiv.org/abs/2006.15417>.

683

684

685

686

687

688

689

690

691

692

693

694

695

696

697

698

699

700

701

702 **A FURTHER RELATED WORKS**

703  
 704  
 705 The theory behind neural networks performing feature decomposition has seen several variations in  
 706 previous years (Bengio et al. 2013), (Locatello et al. 2019). Even prior to studies of latent space  
 707 decomposition, studies found interpretable axes in concept embeddings such as language vectors  
 708 (Mikolov et al. 2013). However, in recent years, we can generally divide discussions of feature  
 709 decomposition into two camps.

710 Empirical papers discuss the decomposition of features for interpretability of specific models or  
 711 domains. Most notable of these are applications of SAEs and similar techniques on LLMs (Bricken  
 712 et al. 2023), (Ameisen et al. 2025), (Gao et al. 2025), which followed the initial results of SAEs  
 713 (Cunningham et al. 2024). These include analyses to extract internal representations of true and false  
 714 statements (Marks et al. 2024), or discovering function vectors by analysing the cumulative impact  
 715 of attention heads (Todd et al. 2024). Similarly, image classification and reconstruction studies have  
 716 proposed new CNN decomposition methods, highlighting segmentations of the input image that led  
 717 to appropriate classification (Ghorbani et al. 2019), (Zhang et al. 2021), (Fel et al. 2023b), (Fel et al.  
 718 2025), (Rao et al. 2024). Particularly, we point out Thasarathan et al. (2025), which suggests that  
 719 text and image features can operate in the same feature space. This suggests that features can bridge  
 720 modalities, and although it doesn't address it independently, the paper invites further investigation  
 721 into the dynamics of cross-modal models. Still other papers propose variations or improvements on  
 722 feature decomposition techniques (Rajamanoharan et al. 2024), (Bussmann et al. 2024) or draw  
 723 comparisons between them (Fel et al. 2023a). Empirical papers typically focus on either the method  
 724 of feature extraction or the application of specifically extracted features towards robustness, safety,  
 725 or interpretability; however, they do not discuss how this feature space was learned by the model or  
 726 generally functions, and, as said before, have left unstructured data domains relatively unexplored.

727 Theoretical papers discuss the structure and formation of the feature space. These papers are fewer  
 728 and further between. We primarily draw on these for our framework of superposition, as these  
 729 papers provided abstracted experiments on feature learning dynamics to build intuition (Elhage et al.  
 730 2022), (Elhage et al. 2023). Rarely, other works have investigated concept learning dynamics  
 731 through accuracy evaluation of individual concepts at each stage of the model (Park et al. 2024).  
 732 However, this body of work has been significantly abstracted away from current state-of-the-art  
 733 models, relying on controlled or toy experiments.

734 In short, there has been a gap in establishing general dynamics of real-world feature spaces. Al-  
 735 though feature decomposition itself has significantly improved with the advent of SAEs, the field  
 736 lacks an equivalent explanation of how these feature spaces are formed. Our work addresses this  
 737 absence though a thorough investigation of a real-world feature space, analyzing the overall trends  
 738 in presence and identity of learned features in addition to their specific function. We also believe the  
 739 gap in generalized feature space research is the result of heavy investment in LLM interpretability  
 740 research. While such studies are clearly highly salient and fruitful, relationships between tokens  
 741 are not explicit, and can be difficult to intuit. In contrast, unstructured and unordered data, while  
 742 difficult to work with, have clear spatial relationships that allow for intuitive interpretation. We  
 743 hope to further discussion regarding properties of the feature space in general, as any insights are  
 744 likely to inform improvements in transparency, robustness, and meta-training.

745 **B ADDITIONAL DISCUSSION**

746  
 747 **B.1 ADDITIONAL SAE VARIATIONS**

748  
 749

n	Relative $\ell_2$ (↓)	Universality (↑)	Dora-VAE Loss (↓)
256	0.518 / 0.366 / 0.194	0.420 / 0.421 / 0.420	0.538 / 0.300 / 0.228
512	0.507 / 0.355 / 0.187	0.293 / 0.297 / 0.295	0.427 / 0.409 / 0.216
1024	0.501 / 0.356 / 0.182	0.206 / 0.208 / 0.209	0.407 / 0.414 / 0.239

750  
 751  
 752  
 753  
 754  
 755 Table 1: SAE Variations, where threshold  $k = 4 / 8 / 16$ .

In addition to the SAE trained in the main paper, we also train several variations on codebook size and threshold  $k$  to establish a general intuition regarding the effects of modifying these hyperparameters. We report the reconstruction loss of the latent, the loss produced by the reconstructed latent passed through the Dora-VAE decoder, and universality. Universality is a measurement of the similarity between features of separately trained SAEs, inspired by Fel et al. (2023a) and Bricken et al. (2023). To measure universality, for each set of hyperparameters, we train 10 identical SAEs on 10-fold subsets of the data. We then use a pairwise Procrustes alignment to align feature vectors between two trained SAEs. Finally, we report the average cosine similarity of vectors among paired models as universality.

We present our results in Table 1. Performance is as expected — as  $k$  increases, the reconstruction is allowed higher fidelity, improving the reconstruction loss of both the SAE and Dora-VAE. Similarly, codebook size improves our metrics, as the dictionary of vectors becomes larger.

The same cannot be said for universality, however; as codebook size increases, features become less universal. We suspect this is due to the many possible tilings of positional encoding — as features are allowed to become more specific, there is a greater variation of feature collections that cover a similar space. In future work, we explore how codebook size affects the patterns of features learned.

## B.2 FURTHER FEATURE ABLATION AND ADDITION INSIGHTS

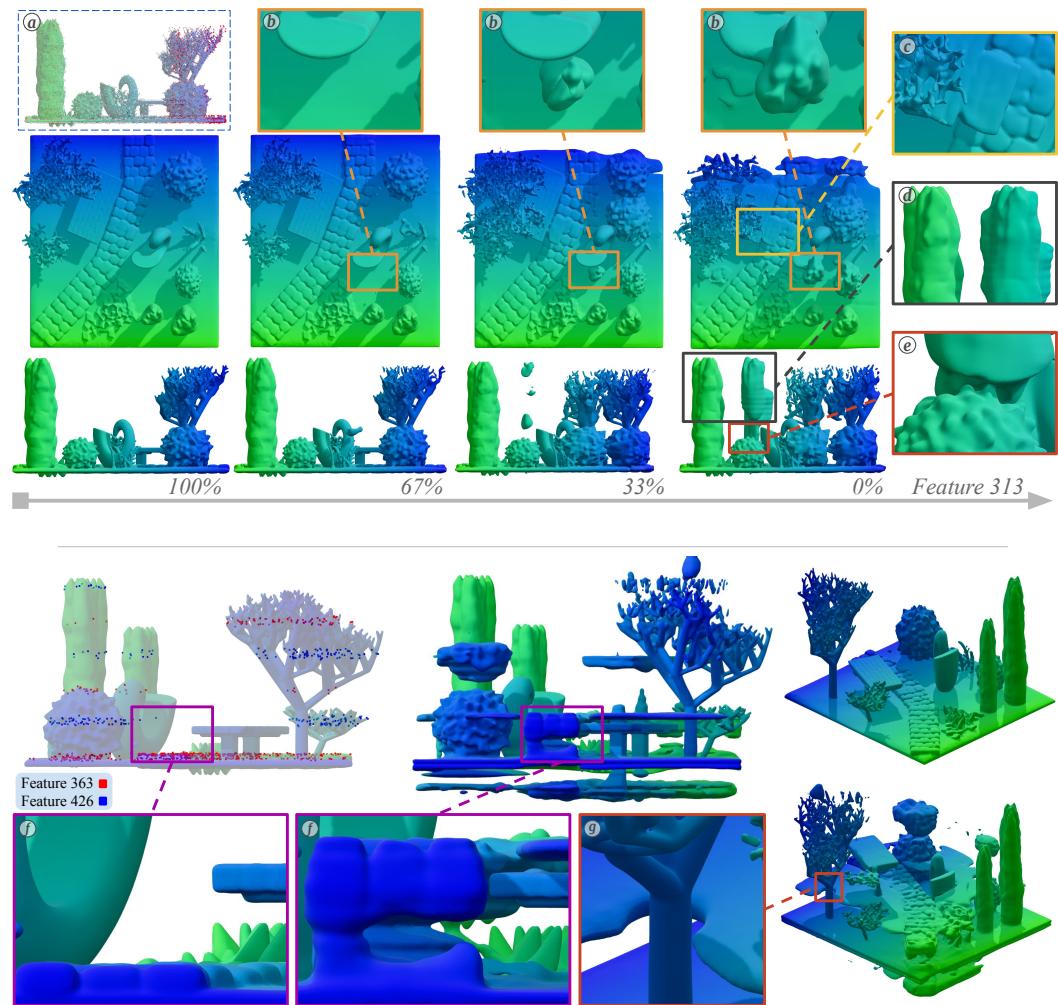


Figure 8: Reprint of Figure 3. *Top*: Example ablation on feature 313. *Bottom*: Example of feature 363 ablation and feature 426 addition.

810 We discuss some interesting properties in Figure 3 here. By gradually removing a single feature, we  
 811 affect the reconstruction by Dora-VAE. In this case, we remove feature 313, which is responsible for  
 812 points along the positive end of the z-axis, as well a small section of points towards the negative end.  
 813 When we remove feature 313, these points move towards the center of the model, forming distinct  
 814 shapes. We frame the discussion around several key properties.

815 As seen in ⑥, the appearance of these shapes is not reflected as a continuous shift in position. In-  
 816 stead, points are spontaneously instantiated in their final position, with more of the shape becoming  
 817 visible over time. This supports our claim that features represent discrete states.

818 In ⑦, we see that points, even after shifts, can merge with their neighbors if they share compatible  
 819 latents. The upper right corner of the bench, originally higher on the z-axis, merges successfully with  
 820 the lower corner of the bench after transposition. This is as opposed to ⑧ and ⑨, where the edges  
 821 of the tree and bush prevent shapes that would render at that location after a shift. We suggest this  
 822 is because latents in the bench have a collaborative relationship with other nearby latents, ensuring  
 823 that nearby points with the same properties will attempt to merge. On the other hand, points on the  
 824 edge of the tree and bush have features for non-occupancy, ensuring that other latents that would  
 825 conflict with the model do not render.

826 ⑩ shows an example of latent redundancy within shapes. We note that, even though a section of  
 827 the bush was removed, the remaining latents were able to reconstruct the local scene with strong  
 828 accuracy. In addition, the new shape rendered by the shifted latents also recreates the bush with  
 829 similar success. We suggest this is due to a set of redundancies within the model decoder to help  
 830 preserve shapes, such that only major changes to latents can affect the model.

831 In the latter figure, we refer to the replacement of feature 363 with feature 426. In short, we substitute  
 832 every presence of feature 363 in the model latents with a presence of equal magnitude in feature 426.  
 833 As both features attend to positions on the y-axis, we see that through this substitution, we displace  
 834 shapes previously located at the region marked in feature 363, and shift them to feature 426. We  
 835 highlight ⑪ here — note that the cobblestone path is now elevated to a region marked by feature  
 836 426, while maintaining its shape. We use this substitution to demonstrate the independence of these  
 837 features, able to affect position while still preserving the majority of the form.

### 839 B.3 LEARNING DYNAMICS’ EFFECT ON DISCRETIZATION

840 One might assume the discretization of features arrives solely from the interference caused by the  
 841 superposition. When a model attempts to represent a large number of features in a small number of  
 842 dimensions, features that appear with low presence are more likely to be construed as interference.  
 843 Thus, the model learns to prefer features that only fire with a significant presence.

844 However, this doesn’t address the effect across all features, nor does it describe the movement of  
 845 parameters that brings about discretization. Notably, models prefer to superimpose features that are  
 846 a) sparse and b) less important. Given that positional information is highly relevant across the entire  
 847 span of inputs, it is unlikely the model would prioritize the information presented by other features  
 848 at the expense of position information.

849 Our framework in Eq. 6 offers a possible explanation to this counter-intuitive behavior, through the  
 850 second term  $\alpha_j \cdot \nabla_{\theta_f} \mathbf{e}_j$ . We see that the signal to the identity of feature  $j$ ,  $\nabla_{\theta_f} \mathbf{e}_j$ , is scaled by the  
 851 presence  $\alpha_j$ . Suppose we had a continuous positional feature  $j$ , where  $\alpha_j(\mathbf{x})$  was higher and lower  
 852 based on whether  $\mathbf{x}$  was closer or further along a designated axis. Across  $\mathbf{x}$  at a variety of positions,  
 853 the identity of feature  $j$  would receive the strongest signal at positions that have a high presence.  
 854 Conversely, during optimization, positions with lower presence of feature  $j$  have to contend with  
 855 other features that have a higher presence, diluting the signal to identity. Over time, we suspect  
 856 that this dynamic drives the identity of  $j$  to solely consider  $\mathbf{x}$  at a specific position with the highest  
 857 presence, localizing the feature. This effect contributes to the discretization of all features, not just  
 858 positional ones.

### 860 B.4 CROSS-CONTAMINATION OF LOSS DIFFERENCE

861 One might object to the supposition that the reconstruction model itself is shifting interference from  
 862 superposition, and instead suggest these shifts reflect a bias in the SAE’s estimations. This objection

864 doesn't argue that the SAE is misestimating based on the importance of the feature (as the SAE does  
 865 not optimize based off the effect of the feature on the final Dora-VAE reconstruction loss); rather, it  
 866 proposes there exists a correlating property with feature importance that consistently increases the  
 867 error of the SAE, causing an offset in its estimation.

868 Given the SAE is a simple two-layer encoder-decoder structure and that encoder and decoder weights  
 869 share a high cosine similarity, if there exists a correlating property, such a feature would affect the  
 870 input with a relatively linear correlation similar to the effect on the estimation observed from the  
 871 output — as this feature increases/decreases, the error of the SAE increases. The two candidates  
 872 with the highest potential for cross-contamination with loss difference in this manner are average  
 873 presence and feature density. We investigate if these two properties drive the relationship between  
 874 loss difference and transition point.

875 To do so, grouping ARCs by feature, we compile loss difference, average presence on active features,  
 876 and feature density as dependent variables, to test their relationship with transition point location.  
 877 For each feature, we perform a linear regression through OLS on these variables, as well as with  
 878 their log forms, to determine which variables scale nonlinearly. We find only loss difference per-  
 879 forms better in log. Then, again for each feature, we perform a partial  $R^2$  analysis to quantify the  
 880 improvement each variable has on predicting transition point. Finally, we take the average  $R^2$  value  
 881 for each of the investigated variables over all features in a defined subset. We repeat this investigation  
 882 four times, considering subsets of features with a number of ablations  $n_{abl} \geq \{500, 1k, 3k, 6k\}$   
 883 ablations, and report our results in Table 2.

$n_{abl}$	# of Features	Log Loss Diff $R^2$	Avg Val $R^2$	Density $R^2$
500	150	0.188±0.104	0.042±0.056	0.025±0.027
1k	128	0.181±0.105	0.046±0.059	0.023±0.025
3k	91	0.183±0.116	0.053±0.064	0.020±0.018
6k	59	0.188±0.119	0.058±0.068	0.017±0.017

891 Table 2: Partial  $R^2$  analysis of variables contributing to transition point.  
 892

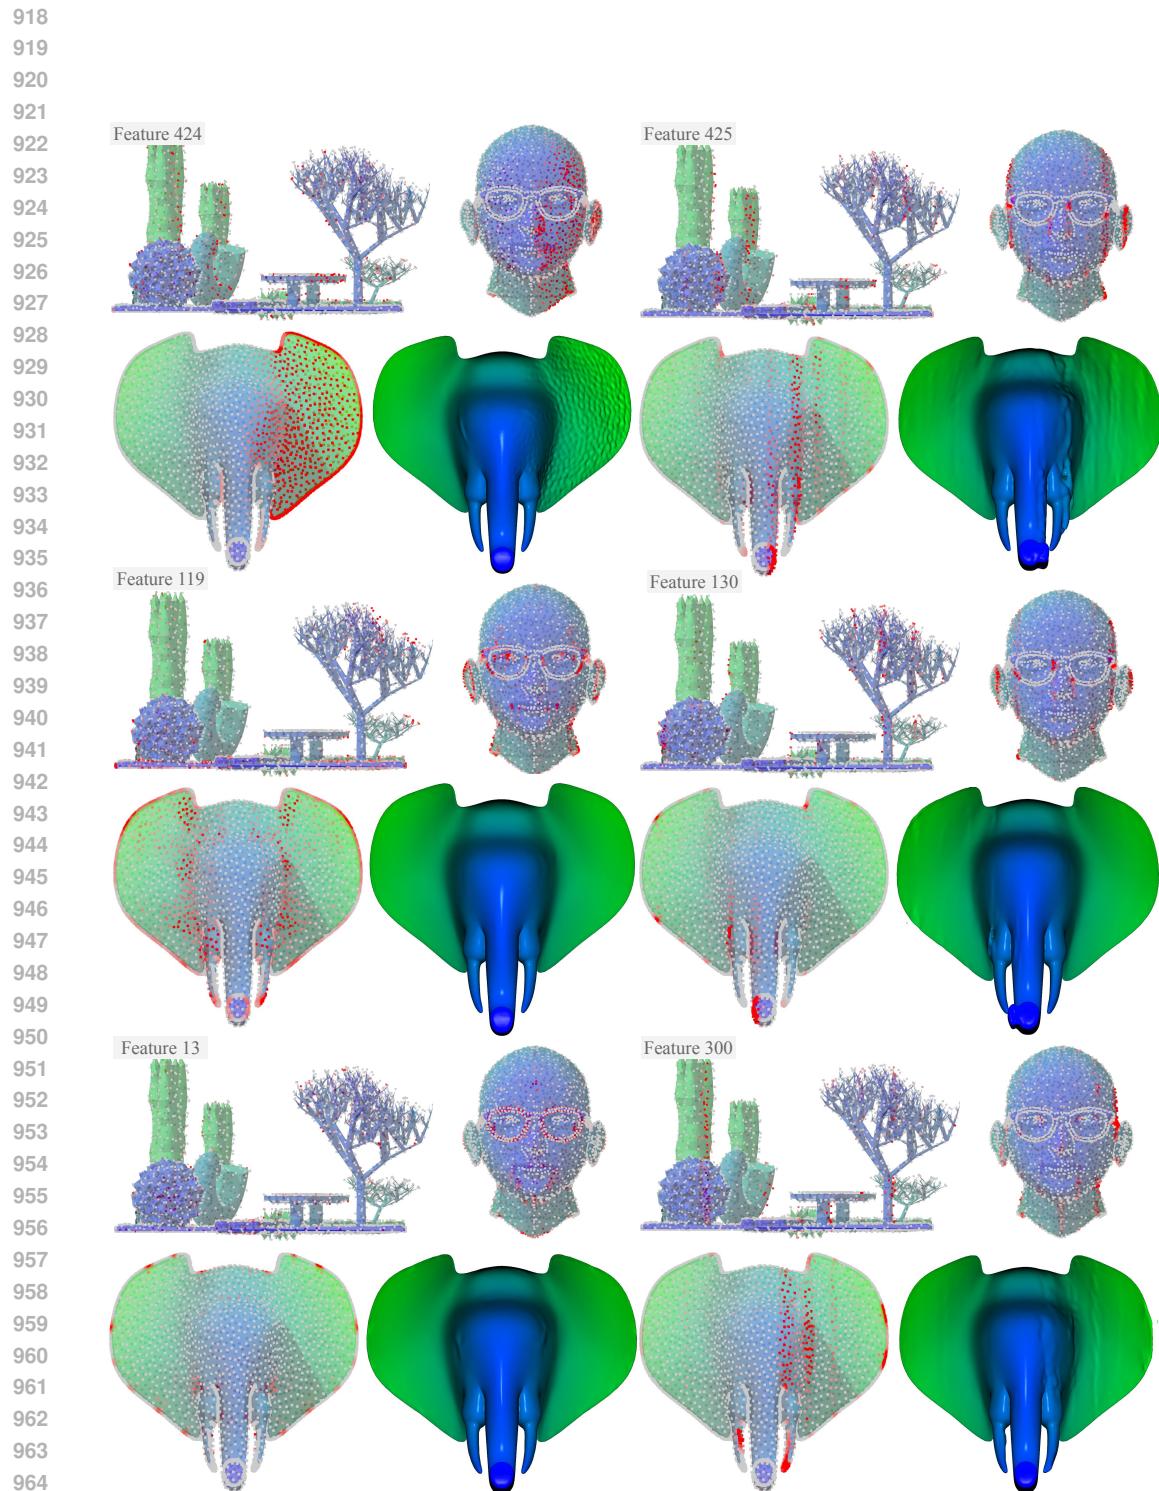
893 Log loss difference consistently contributes to transition point estimation more than other variables.  
 894 This suggests that, rather than either average presence or feature density driving the location of the  
 895 transition point, both variables' effect are likely the result of the correlation with loss difference.  
 896 Thus, the offset in transition points is not caused by an artifact of the SAE.

897 This distinction is important. As opposed to average presence or feature density, neither the dif-  
 898 ference in loss caused by moving the latent in a chosen direction nor the location of the transition  
 899 point is apparent from examining the latent individually. This suggests the model has a complex,  
 900 non-linear relationship with both feature importance and relative transition point, and the correlation  
 901 between these two properties is highly relevant to model behavior.

## 904 C FURTHER FEATURE EXAMINATIONS

905 We demonstrate further examinations of individual features here. Each feature, randomly selected,  
 906 is highlighted across three objects, and ablated for a single object. We also give our qualitative  
 907 impression of each feature's purpose. Note that some ablations have little or no effect, due to the  
 908 redundancy discussed in Appendix B.2, section ④.  
 909

910  
 911  
 912  
 913  
 914  
 915  
 916  
 917



966 Figure 9: *Feature 424*: Smoothness of right-facing surface. *Feature 425*: x-axis encoding,  
967 right side. *Feature 119*: Unknown. *Feature 130*: Left-facing regional smoothness. *Feature*  
968 *13*: Edge encoding. *Feature 300*: x-axis encoding, right side.

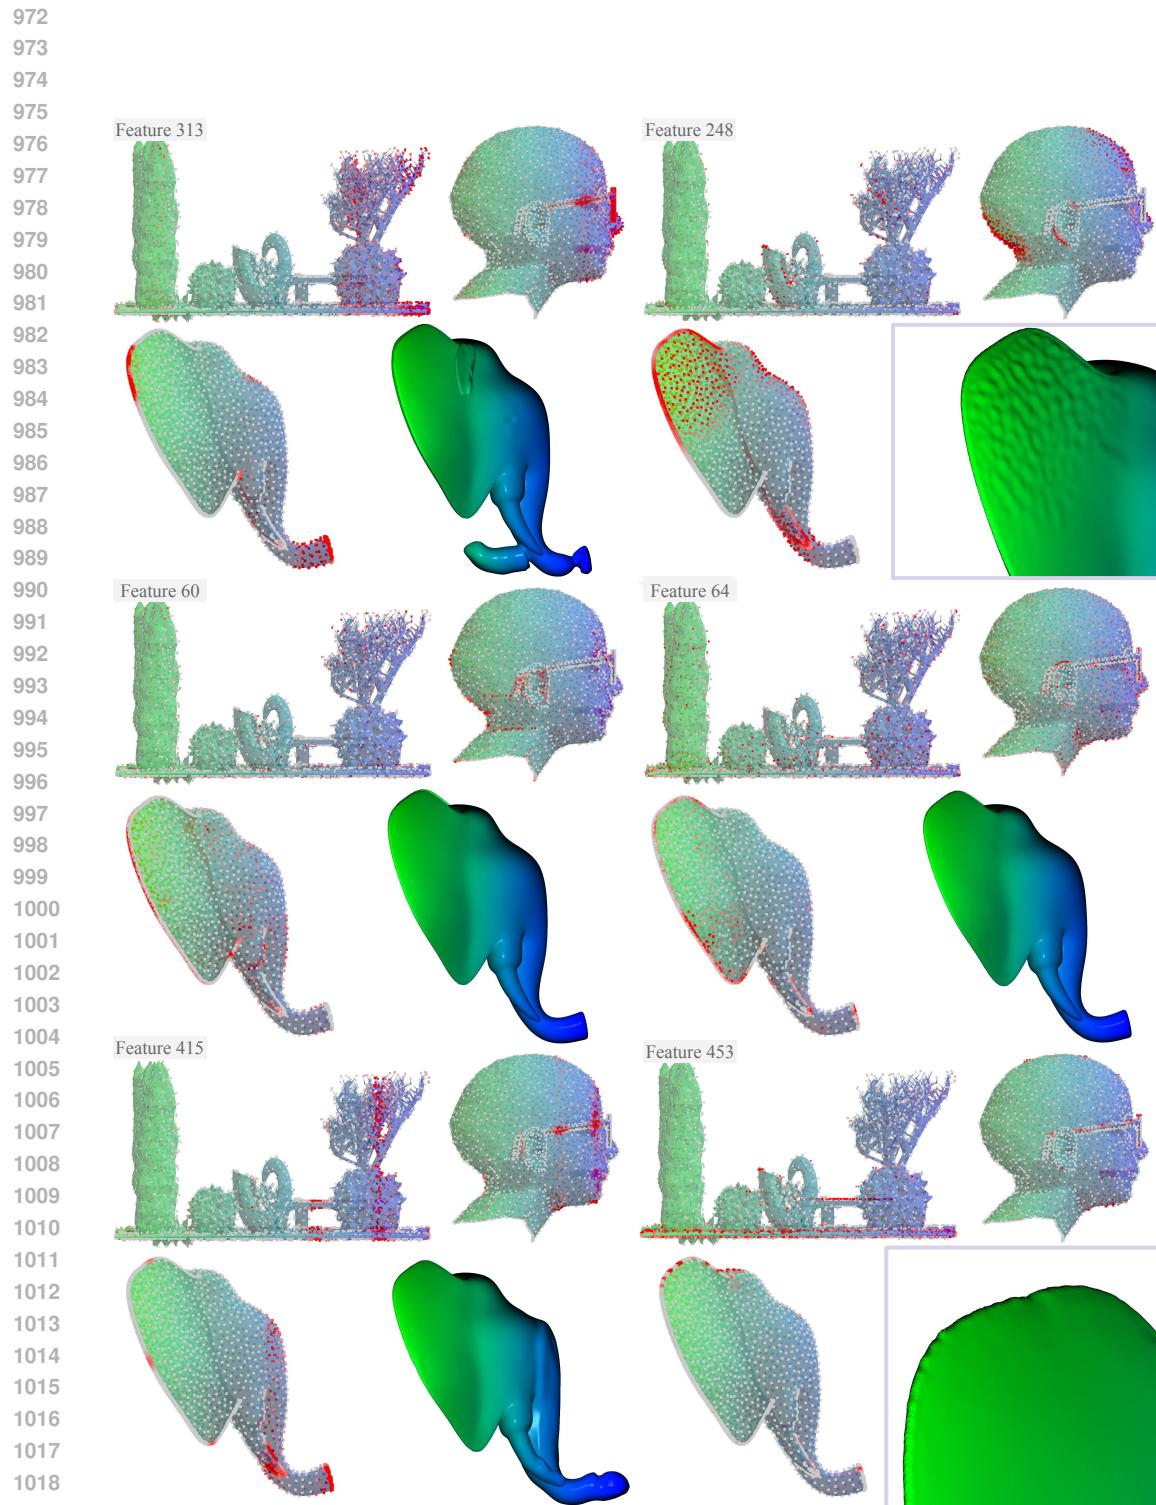


Figure 10: *Feature 313*: z-axis encoding. *Feature 248*: Smoothness of upright-facing surface. *Feature 60*: Unknown. *Feature 64*: Unknown. *Feature 415*: z-axis encoding. *Feature 453*: Upper edge encoding.

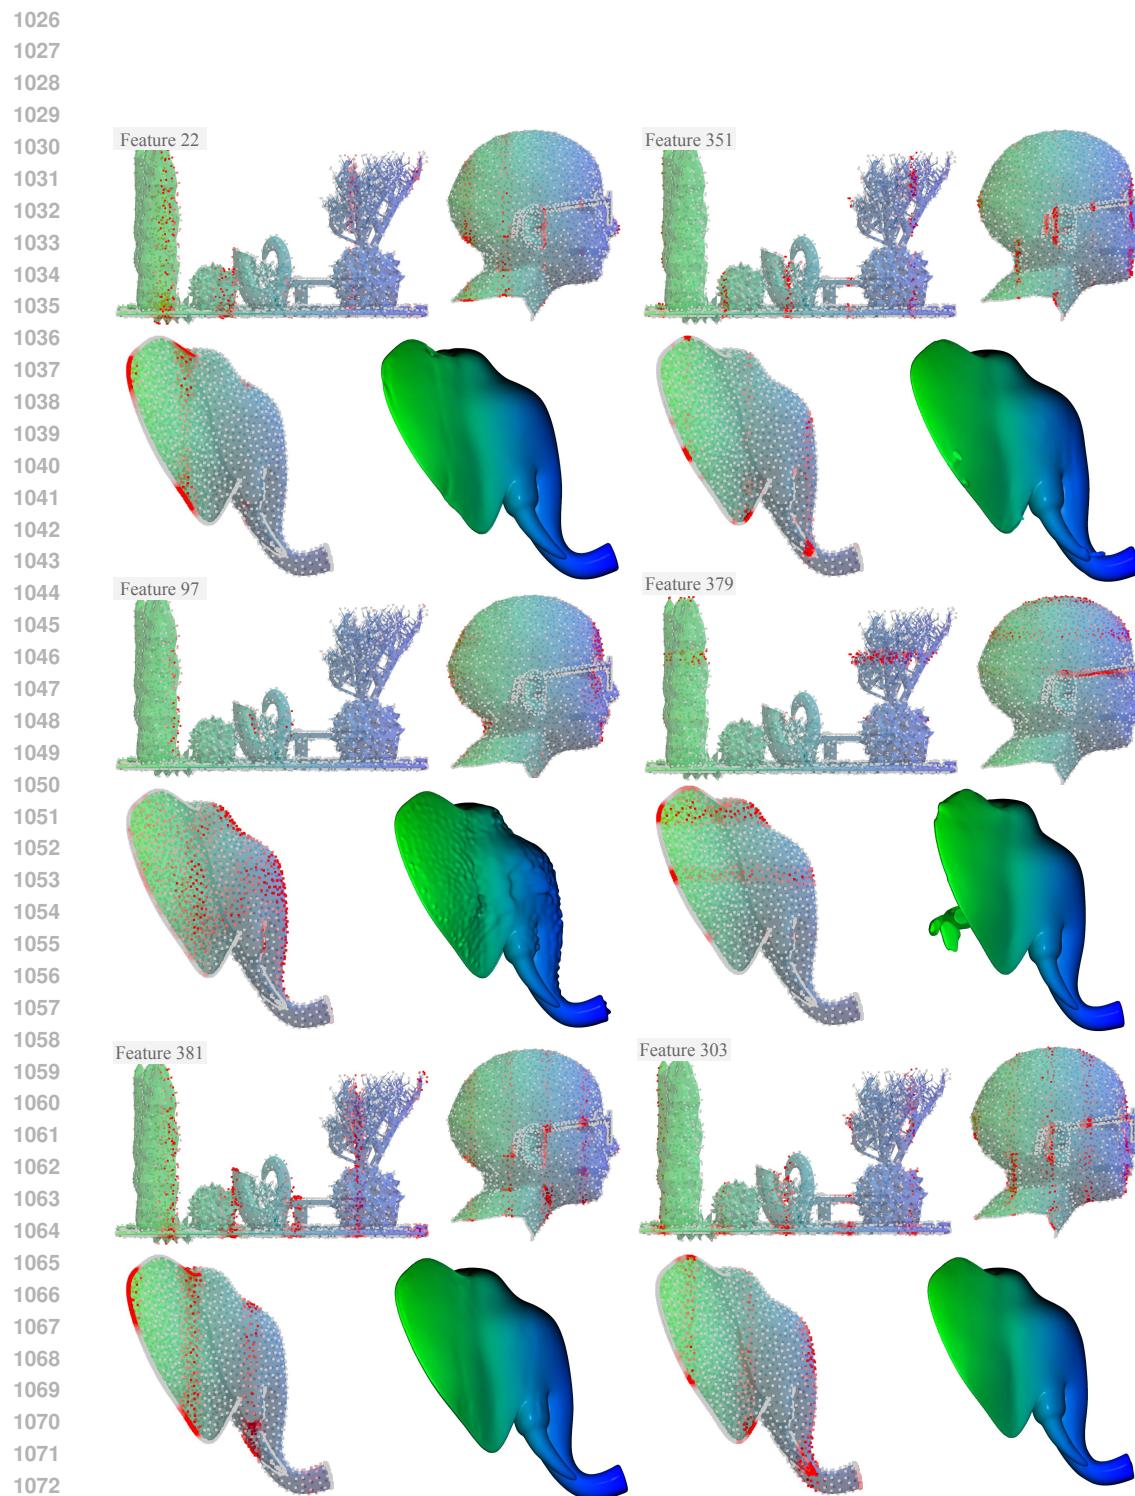


Figure 11: *Feature 22*: z-axis encoding. *Feature 351*: z-axis encoding. *Feature 97*: Right-facing smoothness. *Feature 379*: y-axis encoding. *Feature 381*: z-axis encoding. *Feature 303*: z-axis encoding.