

# SIMULTANEOUS LEARNING FROM BULK AND SINGLE-CELL EXPRESSION DATA WITH PERCEIVER-BASED MODELS

**Anonymous authors**

Paper under double-blind review

## ABSTRACT

In recent years, genomic experiments have generated vast amounts of expression data, including bulk RNA-seq and single-cell RNA-seq datasets, spanning both observational and perturbational studies. However, existing models have not fully leveraged this diverse data landscape, focusing on modeling either bulk or single-cell expression data. We present Funomics T0, the first foundation model that can simultaneously learn from bulk RNA-seq and single-cell RNA-seq datasets from observational and perturbational studies. The proposed Perceiver-based model produces latent representations of the expression data that can be further used for various downstream tasks. We evaluate our model on perturbation prediction and tissue annotation tasks, using a comprehensive benchmark suite and demonstrating strong performance across metrics, with Funomics T0 outperforming the State model on multiple perturbation metrics.

## 1 INTRODUCTION

Next generation sequencing genomic platforms have resulted in the availability of vast amounts of expression data across both bulk and single-cell RNA sequencing modalities. Bulk RNA-seq captures averaged gene expression profiles across heterogeneous mixtures of cells in a tissue or sample, providing an overview of transcriptional activity at the population level, at a small cost. Single-cell RNA-seq, in contrast, quantifies gene expression at the resolution of individual cells, enabling the study of cellular heterogeneity and rare cell populations. These data can originate from observational studies, where native biological variation is explored, or from perturbational experiments designed to measure responses to genetic perturbations, chemical compounds, or other interventions.

The data can be used for various downstream tasks crucial for medical diagnosis and drug discovery, including cell type classification, survival analysis (Wissel et al., 2025), and prediction of perturbation responses (Adduri et al., 2025; Bunne et al., 2023). While these problems can be approached by using small task-specific models, recent advances in foundation models and language modeling have shown that models that can learn from all available data can achieve better performance on various tasks (Brown et al., 2020).

This motivates the recent development of foundation models for transcriptomics that can learn from the vast amount of expression data. This includes foundation models for single-cell RNA-seq data resembling language models, such as scGPT (Cui et al., 2024), scBERT (Yang et al., 2022); models for regulatory inference and gene prioritization (Theodoris et al., 2023; Kalfon et al., 2025); post-perturbation response prediction for genetic or chemical treatments (Adduri et al., 2025; Yang et al., 2024; Wang et al., 2024; Littman et al., 2025); and representation learning for retrieval, zero-shot annotation, or dataset alignment (Rosen et al., 2024; Heimberg et al., 2025). Foundation models for single-cell transcriptomics are primarily based on the transformer architecture (Cui et al., 2024; Yang et al., 2022; Theodoris et al., 2023; Kalfon et al., 2025; Hao et al., 2024; Zeng et al., 2025; Rosen et al., 2024; Yang et al., 2024; Tejada-Lapuerta et al., 2025) with several exceptions based on dense vector or graph representation of the gene expression data (Heimberg et al., 2025).

However, while bulk RNA-seq data is often used for survival analysis (Wissel et al., 2025) and single-cell RNA-seq data is used for cell type classification and perturbation prediction (Adduri et al.,

2025; Bunne et al., 2023), there is no model that can use all available expression data for solving various downstream tasks. This is the gap that we aim to fill with Funomics T0.

## 1.1 CONTRIBUTIONS

We present Funomics T0, a foundation model for transcriptomic modeling that can simultaneously learn from both bulk RNA-seq and single-cell RNA-seq datasets from observational and perturbational studies. Our model learns latent representations of the expression data which can be further used for various downstream tasks. We demonstrate the capabilities of our model on perturbation prediction tasks and tissue annotation tasks.

## 2 METHODS

In this section, we outline the key architectural choices and algorithms behind the Funomics T0 model, which similarly to prior work (Connell et al., 2022; Litman et al., 2025), is based on the Perceiver architecture (Jaegle et al., 2021b;a) and employs the set transformer (Lee et al., 2019) adaptation of cross-attention models widely used in natural language processing following the T5 architecture (Raffel et al., 2023; Zhang et al., 2025a).

The gene expression profiles of either single-cell or bulk RNA-seq are high-dimensional vectors  $X \in \mathbb{R}^G$ , for  $G \approx 19,000$  protein-coding genes. Despite the high dimensionality, the profiles contain substantial biological redundancy due to the functional structure, such as co-regulated pathways, transcriptional programs, and correlated expression patterns (Baek et al., 2025).

If most biologically meaningful variation is concentrated in lower-dimensional expression programs, then a compact representation can faithfully capture underlying structure, while removing the random noise inherent in experimental measurements. To realize this, Funomics T0 introduces a set of  $N_{\text{latent}} \ll G$  learnable latent tokens that serve as a compressed, information-aggregating representation.

This design follows the Perceiver architecture (Jaegle et al., 2021b), in which the model backbone (see Figure 1) consists of a condition encoder and a latent decoder with cross-attention. Gene expression inputs are treated as unordered sets, without positional encodings, reflecting the fact that gene order has no biological meaning. Condition information (*e.g.* perturbation identity, dosage, or metadata) is embedded and processed through a bidirectional self-attention encoder, allowing interactions among condition descriptors. The decoder applies self-attention over the latent tokens, followed by cross-attention to both gene embeddings and condition embeddings, producing compressed latent states  $L \in \mathbb{R}^{N_{\text{latent}} \times d}$ . These latent states are then projected back to gene space via projection cross-attention, producing gene-level hidden states  $H_{\text{genes}} \in \mathbb{R}^{G \times d}$  as visible in Figure 2.

We note that compared to models based on the classical transformer architecture of Vaswani et al. (2017), the complexity of processing the expression representation is reduced from  $\mathcal{O}(G^2)$  operations for naive self-attention to  $\mathcal{O}(G \cdot N_{\text{latent}})$  for the cross-attention based approach. The latent array thus captures structured biological signals while enabling efficient training and inference on full-transcriptome data.

The gene-level hidden states processed by the backbone serve as a shared representation used by lightweight task-specific heads. Different tasks attach their own projection and decoding layers, enabling Funomics T0 to learn from different tasks simultaneously. The detailed forward pass algorithm is provided in Appendix A.

### 2.1 BACKBONE ARCHITECTURE

The backbone of the Funomics T0 model is composed of two primary modules: a condition encoder and a decoder utilizing cross-attention iteratively enhancing the representation of the latent tokens to capture the essential information from the high-dimensional gene expression profile through cross-attention. This architecture employs a set transformer approach (Lee et al., 2019). As a result,

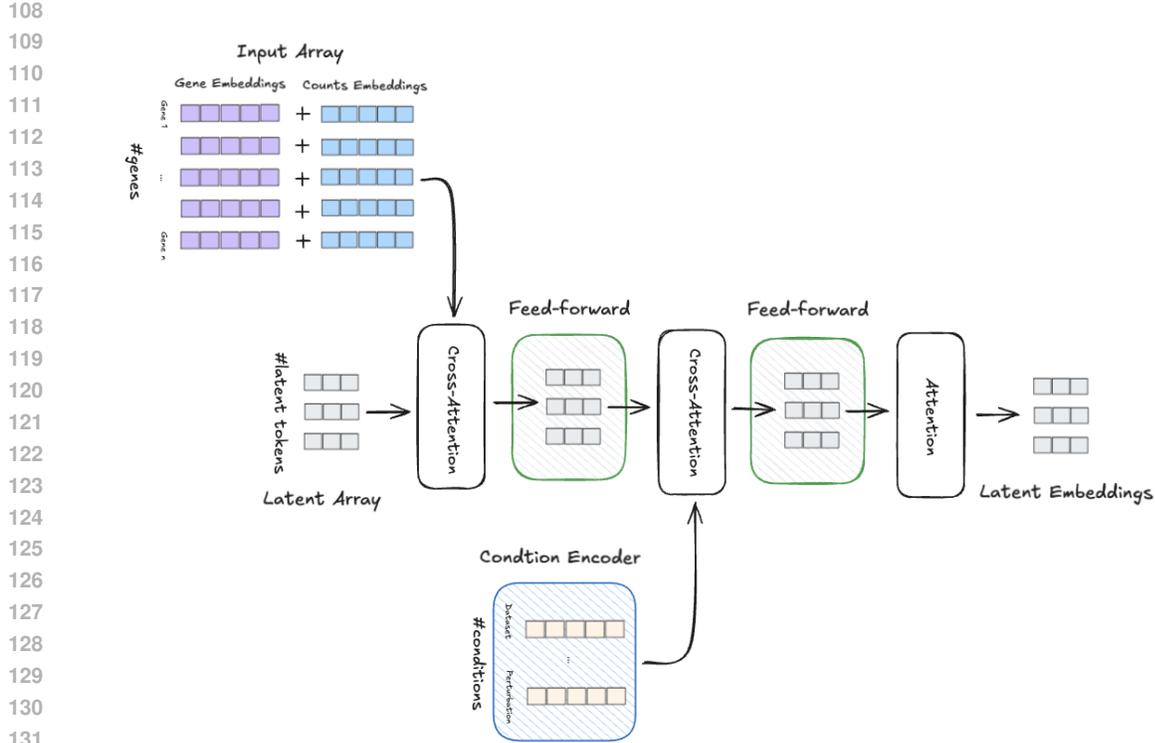


Figure 1: Funomics T0 backbone architecture. The model processes gene expression data through embedding layers, encodes conditions, and uses a decoder with cross-attention to produce compressed latent representations.

the model is inherently permutation-invariant, which aligns well with gene expression data where the ordering of genes is arbitrary. A simplified representation of the backbone architecture is shown in Figure 1.

The condition encoder processes condition embeddings  $E_{\text{cond}} \in \mathbb{R}^{N_{\text{cond}} \times d}$  using bidirectional self-attention, allowing each condition token to attend to all other condition tokens. This enables the encoder to capture relationships between different condition components (e.g., perturbation target gene, dosage, task type) and produce encoded condition representations  $H_{\text{cond}} \in \mathbb{R}^{N_{\text{cond}} \times d}$ .

The decoder operates on a fixed set of  $N_{\text{latent}}$  decoder input embeddings, where  $N_{\text{latent}}$  is chosen to be much smaller than  $G$  ( $N_{\text{latent}} \ll G$ ), achieving significant dimensionality reduction. These decoder inputs are obtained by embedding a fixed vocabulary of  $N_{\text{latent}}$  tokens through learned embedding lookups. The decoder uses self-attention among the latent tokens to enable information exchange, followed by cross-attention to both gene embeddings and encoded condition embeddings. Specifically, gene embeddings  $E_{\text{gene}} \in \mathbb{R}^{G \times d}$  obtained through the gene token embedding module (see H.1) and counts embeddings  $E_{\text{counts}} \in \mathbb{R}^{G \times d}$  through the count embedding module (see H.2) are combined as  $E_{\text{input}} = E_{\text{gene}} + E_{\text{counts}}$  and serve as input to the decoder’s cross-attention mechanism. The decoder also receives encoded condition embeddings  $H_{\text{cond}}$  through encoder-decoder cross-attention, allowing the model to condition its predictions on the specific experimental context.

The decoder produces compressed latent states  $L \in \mathbb{R}^{N_{\text{latent}} \times d}$ , which capture the essential information from the high-dimensional gene expression profile while remaining computationally tractable. These latent states are then projected back to gene space via projection cross-attention, producing gene-level hidden states  $H_{\text{genes}} \in \mathbb{R}^{G \times d}$  that serve as input to task-specific heads. This projection step enables task heads to operate on full-gene representations while maintaining the computational efficiency of the compressed latent processing. The dimensionality reduction achieved by the decoder

162  
163  
164  
165  
166  
167  
168  
169  
170  
171  
172  
173  
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

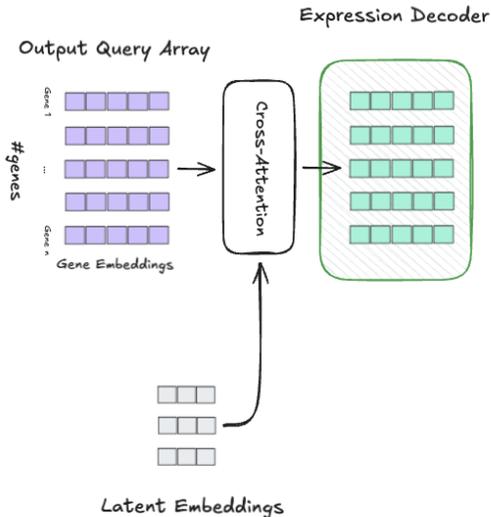


Figure 2: Funomics T0 gene expression decoder. The model uses a linear projection to produce gene-level expression values  $\hat{X} \in \mathbb{R}^{B \times G}$  from gene-level hidden states  $H_{\text{genes}} \in \mathbb{R}^{G \times d}$ .

is crucial for handling the large number of genes ( $G \approx 19,000$ ) efficiently. Architectural details of the backbone components are provided in Appendix E.

### 2.1.1 EMPLOYED EMBEDDINGS

Given expression counts  $X \in \mathbb{R}^{B \times G}$  for a batch of  $B$  cells with  $G$  genes, we construct gene embeddings and counts embeddings independently. Gene embeddings  $E_{\text{gene}} \in \mathbb{R}^{G \times d}$  are learnable token embeddings that provide unique representations for each gene, where  $d$  is the embedding dimension. counts embeddings  $E_{\text{counts}} \in \mathbb{R}^{G \times d}$  encode expression values using either a binning approach or an MLP encoder. These embeddings are combined as  $E_{\text{input}} = E_{\text{gene}} + E_{\text{counts}}$  to form the input representation for the decoder.

Experimental conditions are encoded as separate embeddings that condition the model’s predictions. Condition embeddings handle two distinct cases: genetic perturbations, where target genes are embedded using the same gene token embedding procedure, and chemical perturbations, where molecular embeddings are combined with dosage and optionally the gene targets of a molecule are concatenated.

Task embeddings encode the task type (e.g., perturbation prediction, masked gene prediction, annotation, drug sensitivity). All condition embeddings are concatenated into a condition sequence  $E_{\text{cond}} \in \mathbb{R}^{N_{\text{cond}} \times d}$ , where  $N_{\text{cond}}$  varies based on the perturbation type and task.

Detailed procedures and mathematical formulations for gene token embeddings, counts embeddings, and condition embeddings (including genetic and chemical perturbation embeddings) are provided in Appendix H.

## 2.2 PREDICTION HEADS

Gene-level hidden states  $H_{\text{genes}} \in \mathbb{R}^{G \times d}$  from the Funomics T0 backbone (obtained via projection cross-attention from compressed latent states) are processed by a collection of modular, task-specific prediction heads, each head adapting the backbone-generated compressed representations to the output requirements of its task.

The model employs specialized heads for different tasks: a gene expression decoder for perturbation prediction (producing  $\hat{X} \in \mathbb{R}^{B \times G}$  from  $H_{\text{genes}}$  via linear projection), classification heads that map aggregated gene states to class probabilities ( $\hat{y} \in [0, 1]^{B \times C}$ ), and regression heads for continuous

216 scores ( $\hat{s} \in \mathbb{R}^B$ ). Multi-task heads can output several predictions jointly. Further architectural details  
 217 are provided with each task in Appendix M.  
 218

## 219 2.3 TASK FORMULATIONS

220  
 221 The Funomics T0 model is designed to handle multiple tasks through a unified architecture. Each task  
 222 is formulated as a specific prediction problem that extends the shared backbone with task-specific  
 223 prediction heads (see subsection 2.2). The following subsections describe the formulation for each  
 224 task category. Losses corresponding to each task (see also Appendix J) are then combined into the  
 225 total loss function, used to train the model.  
 226

### 227 2.3.1 PERTURBATION PREDICTIONS

228 In the genetic perturbation prediction task, the model predicts perturbed cell expression profiles given  
 229 control cell counts and genetic perturbation conditions. For a batch of  $B$  cell groups, each containing  
 230  $N$  cells, the model receives control counts  $X_{\text{ctrl}} \in \mathbb{R}^{B \times G}$  and produces predictions  $\hat{X} \in \mathbb{R}^{B \times G}$  for  
 231 the perturbed state, where  $G$  is the number of genes. The target perturbed counts are  $X_{\text{pert}} \in \mathbb{R}^{B \times G}$ .  
 232 Predictions are optionally normalized and can be computed in residual mode by adding control means  
 233 or average perturbation effects. Detailed task description, loss function and data details are provided  
 234 in subsection I.1.1.

235 In the chemical perturbation prediction task, the model predicts perturbed cell expression profiles  
 236 given control cell counts and chemical perturbation conditions (including compound identity and  
 237 dosage). The formulation is similar to genetic perturbations, but condition embeddings incorporate  
 238 molecular representations and dosage information. Detailed task description, loss function and data  
 239 details are provided in subsection I.1.2.  
 240

### 241 2.3.2 UNMASKING PERTURBATION PREDICTION

242 In the unmasking perturbation prediction task, a subset of genes is masked during training, and  
 243 the model predicts expression values for the masked genes. The model receives control counts  
 244  $X_{\text{ctrl}} \in \mathbb{R}^{B \times G}$  with some genes masked, and produces predictions  $\hat{X} \in \mathbb{R}^{B \times G}$  for all genes, where  
 245  $B$  is the batch size and  $G$  is the number of genes.  
 246

247 We perform masking directly on the input count matrix  $X_{\text{ctrl}} \in \mathbb{R}^{B \times G}$ : a binary mask  $M \in$   
 248  $\{0, 1\}^{B \times G}$  is defined where  $M_{b,g} = 1$  indicates that the count value for gene  $g$  in batch  $b$  is  
 249 masked (not observed), and  $M_{b,g} = 0$  indicates it is observed. For each masked entry in  $X_{\text{ctrl}}$ , the  
 250 corresponding value is replaced during count embedding with a special mask token embedding. The  
 251 model is then trained to predict the original observed count value at every masked position.  
 252

### 253 2.3.3 ANNOTATION TASKS

254 In the annotation tasks, the model predicts discrete labels (such as cell type or tissue of origin) from  
 255 gene expression profiles. The model receives expression counts  $X \in \mathbb{R}^{B \times G}$  and produces predictions  
 256  $\hat{y} \in \mathbb{R}^{B \times C}$ , where  $C$  is the number of cell types. Detailed task description, loss function, and data  
 257 details are provided in subsection I.3.  
 258

## 259 3 EXPERIMENTAL RESULTS

### 260 3.1 PERTURBATION PREDICTION

261  
 262 Perturbation prediction is central to drug discovery applications, as it enables *in silico* screening of  
 263 candidate compounds and genetic interventions before costly experimental validation (Baek et al.,  
 264 2025; Csendes et al., 2025). The task involves predicting cellular gene expression responses following  
 265 genetic or chemical perturbations, given control cell expression profiles and perturbation conditions.  
 266 This capability is essential for understanding drug mechanisms of action, identifying off-target effects,  
 267 and prioritizing therapeutic candidates (Wu et al., 2025).  
 268

269 Formally, given control expression profiles  $X_{\text{ctrl}} \in \mathbb{R}^{B \times G}$  for a batch of  $B$  cells and  $G$  genes, along  
 with condition embeddings  $\mathcal{C} \in \mathbb{R}^{B \times N_{\text{cond}} \times d_{\text{cond}}}$ , the model predicts perturbed expression profiles

Table 1: Detailed performance metrics on genetic perturbation prediction.

metric	preprocessing	HEPG2 (Nadig et al., 2025)				K562 (Replogle et al., 2022)			
		State	Funomics T0	mean pert	Bound*	State	Funomics T0	mean pert	Bound*
Cosine Sim. $\uparrow$	train pert centered	0.33	0.41	0.00	0.34	0.21	0.28	0.00	0.34
MSE $\downarrow$	-	0.02	0.03	0.02	0.03	0.02	0.03	0.01	0.01
WMSE $\downarrow$	-	0.02	0.03	0.04	0.02	0.01	0.02	0.02	0.01
Rank.T (Cosine) $\downarrow$	train pert centered	0.22	0.06	1.00	0.11	0.25	0.09	1.00	0.11
Rank.T (L2) $\downarrow$	-	0.26	0.08	0.50	0.25	0.30	0.16	0.50	0.28
Rank (Cosine) $\downarrow$	train pert centered	0.22	0.05	1.00	0.11	0.20	0.05	1.00	0.12
Rank (L2) $\downarrow$	-	0.22	0.07	0.87	0.13	0.22	0.16	0.79	0.10
SMD (Cosine) $\downarrow$	train pert centered	155.19	10.41	11.48	6.60	211.43	14.06	7.32	5.10
SMD (L2) $\downarrow$	-	1562.68	57.09	449.56	55.20	1417.55	44.11	326.49	44.90

\*Bound is the better of technical duplicate and interpolated duplicate baselines.

$\hat{X}_{\text{pert}} \in \mathbb{R}^{B \times G}$ . The goal is to match the distribution of predicted responses,  $\mathbb{P}(\hat{X}_{\text{pert}} | X_{\text{ctrl}}, \mathcal{C})$ , to the true distribution of perturbed responses:  $\mathbb{P}(X_{\text{pert}} | X_{\text{ctrl}}, \mathcal{C})$ , capturing both the magnitude and direction of expression changes across all genes.

**Metrics** To evaluate model performance, we report a compact, complementary metric set that separates magnitude accuracy from directional agreement. We include Cosine similarity (direction), MSE and WMSE (absolute values), rank metrics for Cosine and L2 distances, and similarity matrix distance (SMD) for Cosine and L2. For cosine-based metrics, we center profiles by the training perturbation mean, consistent with prior work Littman et al. (2025). Detailed metric definitions are provided in Appendix K.

**Baselines.** We compare against three baselines (see subsection L.1): (1) **State model** (Adduri et al., 2025), (2) **post-perturbed mean** (see subsection L.1.1), which predicts the mean of perturbed expression in the training dataset for each perturbation following Littman et al. (2025). The post-perturbed mean baseline values were calculated on pseudobulks, without accounting for the number of samples per perturbation which was not present in the pseudobulk data; and (3) **interpolated duplicate baseline** (see subsection L.1.3), which addresses limitations of technical duplicates for weak perturbations (Miller et al., 2025).

We present performance on genetic perturbations (HEPG2 (Nadig et al., 2025) and K562 (Replogle et al., 2022)) in Table 1 and chemical perturbations (Tahoe (Zhang et al., 2025b)) in Table 2. We follow the same splitting strategy as described in the State paper (Adduri et al., 2025) and defer additional dataset results to Appendix D.

### 3.1.1 GENETIC PERTURBATIONS

Overall, Funomics T0 performs better than State on both genetic perturbation datasets on most metrics, with clear gains on cosine-based similarity, rank metrics, and SMD. These gains are especially notable given that Funomics T0 was trained on a broader, more diverse set of datasets, yet still generalizes strongly to these benchmarks. It is slightly lower on absolute-value metrics (MSE/WMSE), though both models perform well and are close to the mean pert baseline; these magnitude errors are less biologically informative. Both models still improve over the mean pert baseline on WMSE, showing WMSE better captures meaningful gains on well-calibrated metrics Mejia et al. (2025). On Nadig HEPG2, we observe results better than the interpolated duplicate bound, which may reflect technical duplicate splits that partition a single-cell population and introduce noise for small populations. We attribute this tradeoff to latent compression, which may soften exact value prediction while preserving the direction of changes that is more biologically relevant.

### 3.1.2 CHEMICAL PERTURBATIONS

On the Tahoe dataset, both models achieve strong performance across metrics, suggesting this benchmark is comparatively easy. State is better on cosine similarity, rank metrics, and MSE/WMSE, indicating more accurate absolute magnitudes and ordering. Funomics T0 is markedly better on SMD, which reflects closer alignment of the predicted perturbation–perturbation similarity structure with the ground truth, even when per-perturbation magnitudes are less calibrated. This suggests

Funomics preserves global relational structure among perturbations, a property that can be useful for downstream grouping and interpretation. The limited separation between models on Tahoe may reflect the dataset’s scale, where single-dataset training already provides enough coverage and reduces the benefit of multi-dataset pretraining.

Table 2: Chemical perturbation prediction results on Tahoe dataset (Zhang et al., 2025b).

metric	preprocessing	State	Funomics	T0
Cosine Sim.	↑ train pert centered	0.78	0.67	
MSE	↓ –	0.01	0.02	
WMSE	↓ –	0.01	0.02	
Rank.T (Cosine)	↓ train pert centered	0.01	0.04	
Rank.T (L2)	↓ –	0.00	0.05	
Rank (Cosine)	↓ train pert centered	0.00	0.01	
Rank (L2)	↓ –	0.01	0.05	
SMD (Cosine)	↓ train pert centered	209.68	65.62	
SMD (L2)	↓ –	803.86	172.57	

### 3.2 TISSUE ANNOTATION

In the tissue annotation task, the model predicts tissue labels from gene expression profiles. For this purpose, we employed the Gene Tissue Expression (GTEx; Lonsdale et al. (2013)) v10 dataset, composed of 19,788 samples from 946 donors, spanning 54 distinct tissue types, profiled using the bulk RNA-sequencing technology. As the model inputs we used the Transcripts Per Million (TPM) normalized expression data. As the number of samples per tissue is quite unbalanced, with a minimum of 11 observations per tissue, and maximum of around eight hundred, we excluded five tissues with less than 100 observations from the validation splits (Kidney-Medulla, Cervix-Ectocervix and -Endocervix, Fallopian Tube, Bladder).

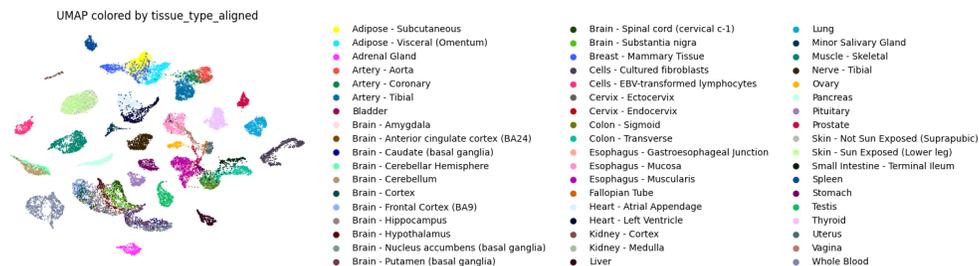


Figure 3: UMAP visualization of GTEx gene expression profiles. Each point represents a bulk sample, colored according to its annotated tissue label.

For the tissue annotation task, we employ two complementary validation splits: `val_purity` (2,426 samples) and `val_tissue` (2,341 samples), with the remaining 14,849 samples used for training. The `val_purity` split is designed to select the hardest samples for model evaluation based on neighborhood purity in PCA space, while the `val_tissue` split provides a tissue-balanced benchmark comparable to prior work. Detailed descriptions of the split algorithms and dataset preprocessing are provided in subsection B.3.

For the tissue annotation task, we tested the performance of the Logistic Regression baseline model on the `val_purity` and `val_tissue` splits. The Logistic Regression model was

378  
379  
380  
381  
382  
383  
384  
385  
386  
387  
388  
389  
390  
391  
392  
393  
394  
395  
396  
397  
398  
399  
400  
401  
402  
403  
404  
405  
406  
407  
408  
409  
410  
411  
412  
413  
414  
415  
416  
417  
418  
419  
420  
421  
422  
423  
424  
425  
426  
427  
428  
429  
430  
431

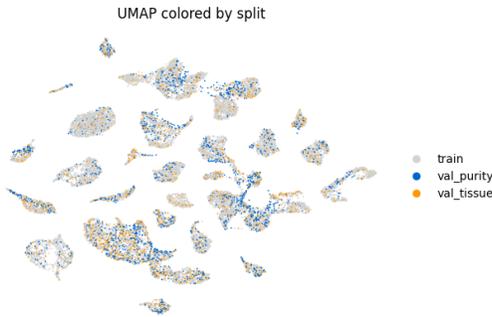


Figure 4: **UMAP visualization of validation splits** GTEX expression data, colored by the splits annotation: `val_purity`, `val_tissue`, `train`.

trained on 5,000 Highly Variable Genes (HVGs) selected using the `scanpy` package, on the `log1p(normalize_total(data))` transformed data and standardized. As both splits were balanced in terms of tissue type, we report Accuracy metric.

Yap et al. (2021) report the F1-score of 96.1% on the split similar to our `val_tissue` split, a value that we report as a reference for our model performance together with the Logistic Regression baseline in Table 3. Additionally, we report the performance of our model trained from scratch (T0) and with pretraining (T0 (pretrained)) on multiple tasks and employing single-cell data. Pretraining yields substantial gains, improving accuracy by over 6 percentage points on `val_purity` and achieving near-perfect performance (99.5%) on `val_tissue`, demonstrating the benefit of large-scale multitask pretraining for downstream tissue classification.

Table 3: **Tissue Annotation Performance Comparison.** Accuracy (in %) of different models on the `val_purity` and `val_tissue` splits of GTEX v10.

Model	<code>val_purity</code> (%)	<code>val_tissue</code> (%)
Funomics T0	86.8	95.7
Funomics T0 (pretrained)	<b>93.2</b>	<b>99.5</b>
Logistic Regression	82.9	98.2

## 4 DISCUSSION AND CONCLUSION

In this work, we present Funomics T0, a foundation model trained to simultaneously learn a shared expression representation across RNA-seq datasets that are typically treated separately: bulk and single-cell, in vitro and patient-derived, and observational and perturbational studies.

The underlying idea is that the main obstacle to a practically useful omics model is the ability to extract stable biological patterns across heterogeneous study designs, parallel biological samples and technical context.

Across perturbation prediction and tissue annotation, we demonstrate its ability to effectively utilize the full spectrum of available transcriptomic data for solving downstream tasks. The latent representations learned by Funomics T0 transfer to task-specific predictors without additional encoders, suggesting that the model captures robust expression programs that are both detectable in single cells, compatible with bulk signals, and present in controlled perturbations and patient data. We see the clearest gains on the smaller genetic perturbation datasets, where multi-dataset training provides stronger generalization. Our results suggest that this unified architecture combined with massive multitask training across heterogeneous data sources can unlock improved modeling of cellular responses, pointing toward a promising direction for deeper understanding of cell processes.

**Limitations and Future Work** While Funomics T0 represents an advancement in foundation models for omics, several limitations remain. Currently, our model focuses on expression data from

bulk and single-cell RNA-seq experiments. Training on heterogeneous RNA-seq gives path to domain imbalance and technical artifacts, and therefore performance can be skewed toward tissues, protocols, or cohorts that are more represented.

Additionally, while we demonstrate the model’s effectiveness on perturbation prediction and tissue annotation tasks, there are numerous other relevant applications to evaluate, including disease classification, drug response prediction, and patient stratification.

Finally, while the latent space is predictive, we aim to interrogate its interpretability and relation to pathways and mechanistic clusters.

Rather than expanding to additional omics modalities, we focus next on expression-based generalization under true biological heterogeneity.

## REFERENCES

- Abhinav K. Adduri, Dhruv Gautam, Beatrice Bevilacqua, Alishba Imran, Rohan Shah, Mohsen Naghipourfar, Noam Teyssier, Rajesh Ilango, Sanjay Nagaraj, Mingze Dong, Chiara Ricci-Tam, Christopher Carpenter, Vishvak Subramanyam, Aidan Winters, Sravya Tirukkovular, Jeremy Sullivan, Brian S. Plosky, Basak Eraslan, Nicholas D. Youngblut, Jure Leskovec, Luke A. Gilbert, Silvana Konermann, Patrick D. Hsu, Alexander Dobin, Dave P. Burke, Hani Goodarzi, and Yusuf H. Roohani. Predicting cellular responses to perturbation across diverse contexts with State. July 2025. doi: 10.1101/2025.06.26.661135. URL <https://www.biorxiv.org/content/10.1101/2025.06.26.661135v2>. ISSN: 2692-8205 Pages: 2025.06.26.661135 Section: New Results.
- Seungbyn Baek, Kyungwoo Song, and Insuk Lee. Single-cell foundation models: bringing artificial intelligence into cell biology. *Experimental & Molecular Medicine*, 57(10):2169–2181, October 2025. ISSN 2092-6413. doi: 10.1038/s12276-025-01547-5. URL <https://www.nature.com/articles/s12276-025-01547-5>.
- Tom Brown, Benjamin Mann, Nick Ryder, Melanie Subbiah, Jared D Kaplan, Prafulla Dhariwal, Arvind Neelakantan, Pranav Shyam, Girish Sastry, Amanda Askell, Sandhini Agarwal, Ariel Herbert-Voss, Gretchen Krueger, Tom Henighan, Rewon Child, Aditya Ramesh, Daniel Ziegler, Jeffrey Wu, Clemens Winter, Chris Hesse, Mark Chen, Eric Sigler, Mateusz Litwin, Scott Gray, Benjamin Chess, Jack Clark, Christopher Berner, Sam McCandlish, Alec Radford, Ilya Sutskever, and Dario Amodei. Language models are few-shot learners. In H. Larochelle, M. Ranzato, R. Hadsell, M.F. Balcan, and H. Lin (eds.), *Advances in Neural Information Processing Systems*, volume 33, pp. 1877–1901. Curran Associates, Inc., 2020. URL [https://proceedings.neurips.cc/paper\\_files/paper/2020/file/1457c0d6bfc4967418bfb8ac142f64a-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2020/file/1457c0d6bfc4967418bfb8ac142f64a-Paper.pdf).
- Charlotte Bunne, Stefan G. Stark, Gabriele Gut, Jacobo Sarabia del Castillo, Mitch Levesque, Kjong-Van Lehmann, Lucas Pelkmans, Andreas Krause, and Gunnar Rätsch. Learning single-cell perturbation responses using neural optimal transport. *Nature Methods*, 20(11):1759–1768, November 2023. ISSN 1548-7105. doi: 10.1038/s41592-023-01969-x. URL <https://www.nature.com/articles/s41592-023-01969-x>.
- William Connell, Umair Khan, and Michael J. Keiser. A single-cell gene expression language model. October 2022. doi: 10.48550/arXiv.2210.14330. URL <http://arxiv.org/abs/2210.14330>. arXiv:2210.14330 [q-bio].
- Gerold Csendes, Gema Sanz, Kristóf Z. Szalay, and Bence Szalai. Benchmarking foundation cell models for post-perturbation RNA-seq prediction. *BMC Genomics*, 26(1):393, April 2025. ISSN 1471-2164. doi: 10.1186/s12864-025-11600-2. URL <https://doi.org/10.1186/s12864-025-11600-2>.
- Haotian Cui, Chloe Wang, Hassaan Maan, Kuan Pang, Fengning Luo, Nan Duan, and Bo Wang. scGPT: toward building a foundation model for single-cell multi-omics using generative AI. *Nature Methods*, 21(8):1470–1480, August 2024. ISSN 1548-7105. doi: 10.1038/s41592-024-02201-0. URL <https://www.nature.com/articles/s41592-024-02201-0>.

- 486 Claudia Feng, Elin Madli Peets, Yan Zhou, Luca Crepaldi, Sunay Usluer, Alistair Dunham, Jana M.  
487 Braunger, Jing Su, Magdalena E. Strauss, Daniele Muraro, Kimberly Ai Xian Cheam, Marc Jan  
488 Bonder, Edgar Garriga Nogales, Sarah Cooper, Andrew Bassett, Steven Leonard, Yong Gu,  
489 Bo Fussing, David Burke, Leopold Parts, Oliver Stegle, and Britta Velten. A genome-scale  
490 single cell CRISPRi map of trans gene regulation across human pluripotent stem cell lines,  
491 November 2024. URL [https://www.biorxiv.org/content/10.1101/2024.11.](https://www.biorxiv.org/content/10.1101/2024.11.28.625833v1)  
492 [28.625833v1](https://www.biorxiv.org/content/10.1101/2024.11.28.625833v1). Pages: 2024.11.28.625833 Section: New Results.
- 493 Minsheng Hao, Jing Gong, Xin Zeng, Chiming Liu, Yucheng Guo, Xingyi Cheng, Taifeng  
494 Wang, Jianzhu Ma, Xuegong Zhang, and Le Song. Large-scale foundation model on  
495 single-cell transcriptomics. *Nature Methods*, 21(8):1481–1491, August 2024. ISSN 1548-  
496 7105. doi: 10.1038/s41592-024-02305-7. URL [https://www.nature.com/articles/](https://www.nature.com/articles/s41592-024-02305-7)  
497 [s41592-024-02305-7](https://www.nature.com/articles/s41592-024-02305-7).
- 498 Graham Heimberg, Tony Kuo, Daryle J. DePianto, Omar Salem, Tobias Heigl, Nathaniel Diamant,  
499 Gabriele Scalia, Tommaso Biancalani, Shannon J. Turley, Jason R. Rock, Héctor Corrada Bravo,  
500 Josh Kaminker, Jason A. Vander Heiden, and Aviv Regev. A cell atlas foundation model for  
501 scalable search of similar human cells. *Nature*, 638(8052):1085–1094, February 2025. ISSN 1476-  
502 4687. doi: 10.1038/s41586-024-08411-y. URL [https://www.nature.com/articles/](https://www.nature.com/articles/s41586-024-08411-y)  
503 [s41586-024-08411-y](https://www.nature.com/articles/s41586-024-08411-y).
- 504 Andrew Jaegle, Sebastian Borgeaud, Jean-Baptiste Alayrac, Carl Doersch, Catalin Ionescu, David  
505 Ding, Skanda Koppula, Daniel Zoran, Andrew Brock, Evan Shelhamer, Olivier J. Henaff, Matthew  
506 Botvinick, Andrew Zisserman, Oriol Vinyals, and Joao Carreira. Perceiver IO: A General Archi-  
507 tecture for Structured Inputs & Outputs. October 2021a. URL [https://openreview.net/](https://openreview.net/forum?id=fILj7WpI-g)  
508 [forum?id=fILj7WpI-g](https://openreview.net/forum?id=fILj7WpI-g).
- 509 Andrew Jaegle, Felix Gimeno, Andy Brock, Oriol Vinyals, Andrew Zisserman, and Joao Car-  
510 reira. Perceiver: General Perception with Iterative Attention. In *Proceedings of the 38th*  
511 *International Conference on Machine Learning*, pp. 4651–4664. PMLR, July 2021b. URL  
512 <https://proceedings.mlr.press/v139/jaegle21a.html>.
- 513 Guillaume Jaume, Paul Doucet, Andrew H Song, Ming Y Lu, Cristina Almagro-Pérez, Sophia J  
514 Wagner, Anurag J Vaidya, Richard J Chen, Drew F K Williamson, Ahrong Kim, and Faisal  
515 Mahmood. HEST-1k: A Dataset for Spatial Transcriptomics and Histology Image Analysis.  
516
- 517 Longda Jiang, Carol Dalgarno, Efthymia Papalexi, Isabella Mascio, Hans-Hermann Wessels, Huiy-  
518 oung Yun, Nika Iremadze, Gila Lithwick-Yanai, Doron Lipson, and Rahul Satija. Systematic  
519 reconstruction of molecular pathway signatures using scalable single-cell perturbation screens,  
520 January 2024. URL <https://zenodo.org/records/10520190>.
- 521 Jérémie Kalfon, Jules Samaran, Gabriel Peyré, and Laura Cantini. scPRINT: pre-training on 50  
522 million cells allows robust gene network predictions. *Nature Communications*, 16(1):3607, April  
523 2025. ISSN 2041-1723. doi: 10.1038/s41467-025-58699-1. URL [https://www.nature.](https://www.nature.com/articles/s41467-025-58699-1)  
524 [com/articles/s41467-025-58699-1](https://www.nature.com/articles/s41467-025-58699-1).
- 525 Juho Lee, Yoonho Lee, Jungtaek Kim, Adam Kosior, Seungjin Choi, and Yee Whye Teh. Set  
526 Transformer: A Framework for Attention-based Permutation-Invariant Neural Networks. In  
527 *Proceedings of the 36th International Conference on Machine Learning*, pp. 3744–3753. PMLR,  
528 May 2019. URL <https://proceedings.mlr.press/v97/lee19d.html>.
- 529 Elon Litman, Tyler Myers, Vinayak Agarwal, Ekansh Mittal, Orion Li, Ashwin Gopinath, and  
530 Timothy Kassis. GeneJepa: A Predictive World Model of the Transcriptome. October 2025. doi:  
531 [10.1101/2025.10.14.682378](https://www.biorxiv.org/content/10.1101/2025.10.14.682378). URL [https://www.biorxiv.org/content/10.1101/](https://www.biorxiv.org/content/10.1101/2025.10.14.682378)  
532 [2025.10.14.682378v1](https://www.biorxiv.org/content/10.1101/2025.10.14.682378). ISSN: 2692-8205 Pages: 2025.10.14.682378 Section: New Results.
- 533 Russell Littman, Jacob Levine, Sepideh Maleki, Yongju Lee, Vladimir Ermakov, Lin Qiu, Alexander  
534 Wu, Kexin Huang, Romain Lopez, Gabriele Scalia, Tommaso Biancalani, David Rich-  
535 mond, Aviv Regev, and Jan-Christian Hütter. Gene-embedding-based prediction and func-  
536 tional evaluation of perturbation expression responses with PRESAGE. June 2025. doi:  
537 [10.1101/2025.06.03.657653](https://www.biorxiv.org/content/10.1101/2025.06.03.657653). URL [https://www.biorxiv.org/content/10.1101/](https://www.biorxiv.org/content/10.1101/2025.06.03.657653)  
538 [2025.06.03.657653v1](https://www.biorxiv.org/content/10.1101/2025.06.03.657653). Pages: 2025.06.03.657653 Section: New Results.
- 539

- 540 Qiyuan Liu, Qirui Zhang, Jinhong Du, Siming Zhao, and Jingshu Wang. Effects of Distance  
541 Metrics and Scaling on the Perturbation Discrimination Score, November 2025. URL <http://arxiv.org/abs/2511.16954>. arXiv:2511.16954 [stat].  
542  
543
- 544 John Lonsdale, Jeffrey Thomas, Mike Salvatore, Rebecca Phillips, Edmund Lo, Saboor Shad, Richard  
545 Hasz, Gary Walters, Fernando Garcia, Nancy Young, Barbara Foster, Mike Moser, Ellen Karasik,  
546 Bryan Gillard, Kimberley Ramsey, Susan Sullivan, Jason Bridge, Harold Magazine, John Syron,  
547 Johnelle Fleming, Laura Siminoff, Heather Traino, Maghboeba Mosavel, Laura Barker, Scott  
548 Jewell, Dan Rohrer, Dan Maxim, Dana Filkins, Philip Harbach, Eddie Cortadillo, Bree Berghuis,  
549 Lisa Turner, Eric Hudson, Kristin Feenstra, Leslie Sobin, James Robb, Phillip Branton, Greg  
550 Korzeniewski, Charles Shive, David Tabor, Liqun Qi, Kevin Groch, Sreenath Nampally, Steve  
551 Buia, Angela Zimmerman, Anna Smith, Robin Burges, Karna Robinson, Kim Valentino, Deborah  
552 Bradbury, Mark Cosentino, Norma Diaz-Mayoral, Mary Kennedy, Theresa Engel, Penelope  
553 Williams, Kenyon Erickson, Kristin Ardlie, Wendy Winckler, Gad Getz, David DeLuca, Daniel  
554 MacArthur, Manolis Kellis, Alexander Thomson, Taylor Young, Ellen Gelfand, Molly Donovan,  
555 Yan Meng, George Grant, Deborah Mash, Yvonne Marcus, Margaret Basile, Jun Liu, Jun Zhu,  
556 Zhidong Tu, Nancy J. Cox, Dan L. Nicolae, Eric R. Gamazon, Hae Kyung Im, Anuar Konkashbaev,  
557 Jonathan Pritchard, Matthew Stevens, Timothée Flutre, Xiaoquan Wen, Emmanouil T. Dermitzakis,  
558 Tuuli Lappalainen, Roderic Guigo, Jean Monlong, Michael Sammeth, Daphne Koller, Alexis Battle,  
559 Sara Mostafavi, Mark McCarthy, Manuel Rivas, Julian Maller, Ivan Rusyn, Andrew Nobel, Fred  
560 Wright, Andrey Shabalina, Mike Feolo, Nataliya Sharopova, Anne Sturcke, Justin Paschal, James M.  
561 Anderson, Elizabeth L. Wilder, Leslie K. Derr, Eric D. Green, Jeffery P. Struwing, Gary Temple,  
562 Simona Volpi, Joy T. Boyer, Elizabeth J. Thomson, Mark S. Guyer, Cathy Ng, Assya Abdallah,  
563 Deborah Colantuoni, Thomas R. Insel, Susan E. Koester, A. Roger Little, Patrick K. Bender,  
564 Thomas Lehner, Yin Yao, Carolyn C. Compton, Jimmie B. Vaught, Sherilyn Sawyer, Nicole C.  
565 Lockhart, Joanne Demchok, and Helen F. Moore. The Genotype-Tissue Expression (GTEx)  
566 project. *Nature Genetics*, 45(6):580–585, June 2013. ISSN 1546-1718. doi: 10.1038/ng.2653.  
567 URL <https://www.nature.com/articles/ng.2653>.
- 568 Gabriel M. Mejia, Henry E. Miller, Francis J. A. Leblanc, Bo Wang, Brendan Swain, and Lucas  
569 Paulo de Lima Camillo. Diversity by Design: Addressing Mode Collapse Improves scRNA-seq  
570 Perturbation Modeling on Well-Calibrated Metrics, June 2025. URL <http://arxiv.org/abs/2506.22641>. arXiv:2506.22641 [q-bio].
- 571 Henry E. Miller, Gabriel M. Mejia, Francis J. A. Leblanc, Bo Wang, Brendan Swain, and Lucas Paulo  
572 de Lima Camillo. Deep Learning-Based Genetic Perturbation Models Do Outperform Uninformative  
573 Baselines on Well-Calibrated Metrics. October 2025. doi: 10.1101/2025.10.20.683304. URL  
574 <https://www.biorxiv.org/content/10.1101/2025.10.20.683304v1>. ISSN:  
575 2692-8205 Pages: 2025.10.20.683304 Section: New Results.
- 576 Pablo Moreno, Silvie Fexova, Nancy George, Jonathan R Manning, Zhichiao Miao, Suhaib Mo-  
577 hammed, Alfonso Muñoz-Pomer, Anja Fullgrabe, Yalan Bi, Natassja Bush, Haider Iqbal, Upendra  
578 Kumbham, Andrey Solovyev, Lingyun Zhao, Ananth Prakash, David García-Seisdedos, Deepti J  
579 Kundu, Shengbo Wang, Mathias Walzer, Laura Clarke, David Osumi-Sutherland, Marcela Karey  
580 Tello-Ruiz, Sunita Kumari, Doreen Ware, Jana Eliasova, Mark J Arends, Martijn C Nawijn, Kerstin  
581 Meyer, Tony Burdett, John Marioni, Sarah Teichmann, Juan Antonio Vizcaíno, Alvis Brazma, and  
582 Irene Papatheodorou. Expression Atlas update: gene and protein expression in multiple species.  
583 *Nucleic Acids Research*, 50(D1):D129–D140, November 2021. ISSN 0305-1048. doi: 10.1093/  
584 nar/gkab1030. URL <https://pmc.ncbi.nlm.nih.gov/articles/PMC8728300/>.
- 585  
586 Ajay Nadig, Joseph M. Replogle, Angela N. Pogson, Mukundh Murthy, Steven A. McCarroll,  
587 Jonathan S. Weissman, Elise B. Robinson, and Luke J. O’Connor. Transcriptome-wide analysis  
588 of differential expression in perturbation atlases. *Nature Genetics*, 57(5):1228–1237, May 2025.  
589 ISSN 1546-1718. doi: 10.1038/s41588-025-02169-3. URL [https://www.nature.com/](https://www.nature.com/articles/s41588-025-02169-3)  
590 [articles/s41588-025-02169-3](https://www.nature.com/articles/s41588-025-02169-3).
- 591 Colin Raffel, Noam Shazeer, Adam Roberts, Katherine Lee, Sharan Narang, Michael Matena,  
592 Yanqi Zhou, Wei Li, and Peter J. Liu. Exploring the Limits of Transfer Learning with a Unified  
593 Text-to-Text Transformer, September 2023. URL <http://arxiv.org/abs/1910.10683>.  
arXiv:1910.10683 [cs].

- 594 Joseph M. Replogle, Reuben A. Saunders, Angela N. Pogson, Jeffrey A. Hussmann, Alexander  
595 Lenail, Alina Guna, Lauren Mascibroda, Eric J. Wagner, Karen Adelman, Gila Lithwick-Yanai,  
596 Nika Iremadze, Florian Oberstrass, Doron Lipson, Jessica L. Bonnar, Marco Jost, Thomas M.  
597 Norman, and Jonathan S. Weissman. Mapping information-rich genotype-phenotype landscapes  
598 with genome-scale Perturb-seq. *Cell*, 185(14):2559–2575.e28, July 2022. ISSN 0092-8674, 1097-  
599 4172. doi: 10.1016/j.cell.2022.05.013. URL [https://www.cell.com/cell/abstract/  
600 S0092-8674\(22\)00597-9](https://www.cell.com/cell/abstract/S0092-8674(22)00597-9).
- 601 Yanay Rosen, Yusuf Roohani, Ayush Agrawal, Leon Samotorčan, Tabula Sapiens Consortium,  
602 Stephen R. Quake, and Jure Leskovec. Universal Cell Embeddings: A Foundation Model for Cell  
603 Biology. October 2024. doi: 10.1101/2023.11.28.568918. URL [https://www.biorxiv.  
604 org/content/10.1101/2023.11.28.568918v2](https://www.biorxiv.org/content/10.1101/2023.11.28.568918v2). Pages: 2023.11.28.568918 Section:  
605 New Results.
- 606 Sanjay R. Srivatsan, José L. McFaline-Figueroa, Vijay Ramani, Lauren Saunders, Junyue Cao,  
607 Jonathan Packer, Hannah A. Pliner, Dana L. Jackson, Riza M. Daza, Lena Christiansen, Fan Zhang,  
608 Frank Steemers, Jay Shendure, and Cole Trapnell. Massively multiplex chemical transcriptomics  
609 at single-cell resolution. *Science*, 367(6473):45–51, January 2020. doi: 10.1126/science.aax6234.  
610 URL <https://www.science.org/doi/10.1126/science.aax6234>.
- 611 Alejandro Tejada-Lapueta, Anna C. Schaar, Robert Gutgesell, Giovanni Palla, Lennard Halle, Mariia  
612 Minaeva, Larsen Vornholz, Leander Dony, Francesca Drummer, Till Richter, Mojtaba Bahrami, and  
613 Fabian J. Theis. Nicheformer: a foundation model for single-cell and spatial omics. *Nature Methods*,  
614 22(12):2525–2538, December 2025. ISSN 1548-7105. doi: 10.1038/s41592-025-02814-z. URL  
615 <https://www.nature.com/articles/s41592-025-02814-z>.
- 616 Christina V. Theodoris, Ling Xiao, Anant Chopra, Mark D. Chaffin, Zeina R. Al Sayed, Matthew C.  
617 Hill, Helene Mantineo, Elizabeth M. Brydon, Zexian Zeng, X. Shirley Liu, and Patrick T. Ellinor.  
618 Transfer learning enables predictions in network biology. *Nature*, 618(7965):616–624, June 2023.  
619 ISSN 1476-4687. doi: 10.1038/s41586-023-06139-9. URL [https://www.nature.com/  
620 articles/s41586-023-06139-9](https://www.nature.com/articles/s41586-023-06139-9).
- 621 Ashish Vaswani, Noam Shazeer, Niki Parmar, Jakob Uszkoreit, Llion Jones, Aidan N Gomez,  
622 Łukasz Kaiser, and Illia Polosukhin. Attention is All you Need. In I. Guyon, U. Von  
623 Luxburg, S. Bengio, H. Wallach, R. Fergus, S. Vishwanathan, and R. Garnett (eds.), *Ad-  
624 vances in Neural Information Processing Systems*, volume 30. Curran Associates, Inc.,  
625 2017. URL [https://proceedings.neurips.cc/paper\\_files/paper/2017/  
626 file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf](https://proceedings.neurips.cc/paper_files/paper/2017/file/3f5ee243547dee91fbd053c1c4a845aa-Paper.pdf).
- 627 Gefei Wang, Tianyu Liu, Jia Zhao, Youshu Cheng, and Hongyu Zhao. Modeling and predicting  
628 single-cell multi-gene perturbation responses with scLAMBDA. December 2024. doi: 10.1101/  
629 2024.12.04.626878. URL [https://www.biorxiv.org/content/10.1101/2024.12.  
630 04.626878v1](https://www.biorxiv.org/content/10.1101/2024.12.04.626878v1). Pages: 2024.12.04.626878 Section: New Results.
- 631 David Wissel, Nikita Janakarajan, Aayush Grover, Enrico Toniato, Maria Rodríguez Martínez,  
632 and Valentina Boeva. SurvBoard: standardized benchmarking for multi-omics cancer survival  
633 models. *Briefings in Bioinformatics*, 26(5):bbaf521, September 2025. ISSN 1477-4054. doi:  
634 10.1093/bib/bbaf521. URL <https://doi.org/10.1093/bib/bbaf521>.
- 635 Yan Wu, Esther Wershof, Sebastian M. Schmon, Marcel Nassar, Błażej Osiński, Ridvan Eksi, Zichao  
636 Yan, Rory Stark, Kun Zhang, and Thore Graepel. PerturBench: Benchmarking Machine Learning  
637 Models for Cellular Perturbation Analysis. October 2025. doi: 10.48550/arXiv.2408.10609. URL  
638 <http://arxiv.org/abs/2408.10609>. arXiv:2408.10609 [cs].
- 639 Fan Yang, Wenchuan Wang, Fang Wang, Yuan Fang, Duyu Tang, Junzhou Huang, Hui Lu, and  
640 Jianhua Yao. scBERT as a large-scale pretrained deep language model for cell type annotation of  
641 single-cell RNA-seq data. *Nature Machine Intelligence*, 4(10):852–866, October 2022. ISSN 2522-  
642 5839. doi: 10.1038/s42256-022-00534-z. URL [https://www.nature.com/articles/  
643 s42256-022-00534-z](https://www.nature.com/articles/s42256-022-00534-z).
- 644 Xiaodong Yang, Guole Liu, Guihai Feng, Dechao Bu, Pengfei Wang, Jie Jiang, Shubai Chen,  
645 Qinqing Yang, Hefan Miao, Yiyang Zhang, Zhenpeng Man, Zhongming Liang, Zichen Wang,  
646  
647

648 Yaning Li, Zheng Li, Yana Liu, Yao Tian, Wenhao Liu, Cong Li, Ao Li, Jingxi Dong, Zhilong  
649 Hu, Chen Fang, Lina Cui, Zixu Deng, Haiping Jiang, Wentao Cui, Jiahao Zhang, Zhaohui Yang,  
650 Handong Li, Xingjian He, Liqun Zhong, Jiaheng Zhou, Zijian Wang, Qingqing Long, Ping Xu,  
651 Xin Li, Hongmei Wang, Zhen Meng, Xuezhi Wang, Yangang Wang, Yong Wang, Shihua Zhang,  
652 Jingtao Guo, Yi Zhao, Yuanchun Zhou, Fei Li, Jing Liu, Yiqiang Chen, Ge Yang, and Xin Li.  
653 GeneCompass: deciphering universal gene regulatory mechanisms with a knowledge-informed  
654 cross-species foundation model. *Cell Research*, 34(12):830–845, December 2024. ISSN 1748-  
655 7838. doi: 10.1038/s41422-024-01034-y. URL [https://www.nature.com/articles/  
656 s41422-024-01034-y](https://www.nature.com/articles/s41422-024-01034-y).

657  
658  
659 Melvyn Yap, Rebecca L. Johnston, Helena Foley, Samuel MacDonald, Olga Kondrashova,  
660 Khoa A. Tran, Katia Nones, Lambros T. Koufariotis, Cameron Bean, John V. Pearson, Ma-  
661 cieJ Trzaskowski, and Nicola Waddell. Verifying explainability of a deep learning tissue clas-  
662 sifier trained on RNA-seq data. *Scientific Reports*, 11(1):2641, January 2021. ISSN 2045-  
663 2322. doi: 10.1038/s41598-021-81773-9. URL [https://www.nature.com/articles/  
664 s41598-021-81773-9](https://www.nature.com/articles/s41598-021-81773-9).

665  
666  
667 Yuansong Zeng, Jiancong Xie, Ningyuan Shangguan, Zhuoyi Wei, Wenbing Li, Yun Su, Shuangyu  
668 Yang, Chengyang Zhang, Jinbo Zhang, Nan Fang, Hongyu Zhang, Yutong Lu, Huiying Zhao,  
669 Jue Fan, Weijiang Yu, and Yuedong Yang. CellFM: a large-scale foundation model pre-trained  
670 on transcriptomics of 100 million human cells. *Nature Communications*, 16(1):4679, May 2025.  
671 ISSN 2041-1723. doi: 10.1038/s41467-025-59926-5. URL [https://www.nature.com/  
672 articles/s41467-025-59926-5](https://www.nature.com/articles/s41467-025-59926-5).

673  
674  
675 Biao Zhang, Fedor Moiseev, Joshua Ainslie, Paul Suganthan, Min Ma, Surya Bhupatiraju, Fede  
676 Lebron, Orhan Firat, Armand Joulin, and Zhe Dong. Encoder-Decoder Gemma: Improving the  
677 Quality-Efficiency Trade-Off via Adaptation, April 2025a. URL [http://arxiv.org/abs/  
678 2504.06225](http://arxiv.org/abs/2504.06225). arXiv:2504.06225 [cs].

679  
680  
681 Jesse Zhang, Airoi A. Ubas, Richard de Borja, Valentine Svensson, Nicole Thomas, Neha Thakar,  
682 Ian Lai, Aidan Winters, Umair Khan, Matthew G. Jones, Vuong Tran, Joseph Pangallo, Efthymia  
683 Papalexi, Ajay Sapre, Hoai Nguyen, Oliver Sanderson, Maria Nigos, Olivia Kaplan, Sarah  
684 Schroeder, Bryan Hariadi, Simone Marrujo, Crina Curca Alec Salvino, Guillermo Gallareta  
685 Olivares, Ryan Koehler, Gary Geiss, Alexander Rosenberg, Charles Roco, Daniele Merico, Nima  
686 Alidoust, Hani Goodarzi, and Johnny Yu. Tahoe-100M: A Giga-Scale Single-Cell Perturbation  
687 Atlas for Context-Dependent Gene Function and Cellular Modeling, February 2025b. URL  
688 <https://www.biorxiv.org/content/10.1101/2025.02.20.639398v1>. Pages:  
689 2025.02.20.639398 Section: New Results.

690  
691  
692 Yiming Zhang, Ting Zhang, Gaoxia Yang, Zhenzhong Pan, Min Tang, Yue Wen, Ping He, Yuan Wang,  
693 and Ran Zhou. PerturbAtlas: a comprehensive atlas of public genetic perturbation bulk RNA-seq  
694 datasets. *Nucleic Acids Research*, 53(D1):D1112–D1119, January 2025c. ISSN 1362-4962. doi:  
695 10.1093/nar/gkae851. URL <https://doi.org/10.1093/nar/gkae851>.

696  
697  
698  
699  
700  
701

## A FULL FORWARD PASS ALGORITHM

The complete forward pass algorithm (algorithm 1) orchestrates three main components: task-specific input processing (algorithm 2), backbone processing (algorithm 3), and task head application (algorithm 4).

---

**Algorithm 1:** Funomics T0 Forward Pass

---

```

702
703
704 1 Function ForwardPass ( $X, \mathcal{C}, \mathcal{T}$ ):
705   Input: Gene expression  $X \in \mathbb{R}^{B \times G}$ , conditions  $\mathcal{C}$ , task  $\mathcal{T}$ 
706   Output: Task outputs and losses
707 2    $X_{\text{proc}}, \mathcal{C}_{\text{proc}} \leftarrow \text{ProcessTaskInputs}(X, \mathcal{C}, \mathcal{T})$ 
708   Process inputs according to task requirements (see algorithm 2)
709 3    $L \leftarrow \text{ProcessBackbone}(X_{\text{proc}}, \mathcal{C}_{\text{proc}})$ 
710   Process through backbone to obtain latent states (see algorithm 3)
711 4    $\mathcal{O}, \mathcal{L} \leftarrow \text{ApplyTaskHead}(L, \mathcal{T})$ 
712   Apply task-specific head to generate outputs (see algorithm 4)
713 5   return  $\mathcal{O}, \mathcal{L}$ 

```

---

**Algorithm 2:** Task-Specific Input Processing

---

```

714
715
716 1 Function ProcessTaskInputs ( $X, \mathcal{C}, \mathcal{T}$ ):
717   Input: Gene expression  $X$ , conditions  $\mathcal{C}$ , task  $\mathcal{T}$ 
718   Output: Processed inputs  $X_{\text{proc}}, \mathcal{C}_{\text{proc}}$ 
719 2   if  $\mathcal{T}$  is masked counts task then
720 3     Apply masking to  $X$  according to task-specific masking strategy
721 4   else
722 5     if  $\mathcal{T}$  is perturbation prediction task then
723 6       Prepare control and perturbed counts
724       Apply augmentation masking to control counts if training
725 7     else
726 8       Use inputs as provided
727 9   Process conditions  $\mathcal{C}$  according to task requirements
728   Extract relevant condition components (dataset, task, perturbation
729   type)
730 10  return  $X_{\text{proc}}, \mathcal{C}_{\text{proc}}$ 

```

---

731  
732  
733  
734  
735  
736  
737  
738  
739  
740  
741  
742  
743  
744  
745  
746  
747  
748  
749  
750  
751  
752  
753  
754  
755

**Algorithm 3:** Funomics T0 Backbone Processing

---

```

756
757
758 1 Function ProcessBackbone ( $X, \mathcal{C}$ ):
759   Input: Gene expression  $X \in \mathbb{R}^{B \times G}$ , conditions  $\mathcal{C}$ 
760   Output: Latent states  $L \in \mathbb{R}^{B \times N_{\text{latent}} \times d}$ 
761 2    $E_{\text{gene}}, E_{\text{counts}}, E_{\text{cond}} \leftarrow \text{EmbedGeneExpressionAndConditions}(X, \mathcal{C})$ 
762   Embed gene expression and conditions (see Appendix H)
763    $E_{\text{gene}} \in \mathbb{R}^{G \times d}$ ,  $E_{\text{counts}} \in \mathbb{R}^{G \times d}$ ,  $E_{\text{cond}} \in \mathbb{R}^{N_{\text{cond}} \times d}$ 
764 3    $E_{\text{input}} \leftarrow E_{\text{gene}} + E_{\text{counts}}$ 
765   Combine gene and counts embeddings:  $E_{\text{input}} \in \mathbb{R}^{G \times d}$ 
766 4    $H_{\text{cond}} \leftarrow \text{EncodeConditions}(E_{\text{cond}})$ 
767   Encode conditions through bidirectional self-attention
768   Output:  $H_{\text{cond}} \in \mathbb{R}^{N_{\text{cond}} \times d}$ 
769 5    $L \leftarrow \text{CompressToLatents}(E_{\text{input}}, H_{\text{cond}})$ 
770   Compress gene data to latent representation via decoder
771   cross-attention
772   Decoder uses self-attention on latent tokens, then cross-attention
773   to  $E_{\text{input}}$  and  $H_{\text{cond}}$ 
774   Output:  $L \in \mathbb{R}^{N_{\text{latent}} \times d}$ 
775 6   return  $L$ 
776
777 7 Function EmbedGeneExpressionAndConditions ( $X, \mathcal{C}$ ):
778 8    $E_{\text{gene}} \leftarrow \text{Embed gene tokens}$ 
779   Lookup gene token embeddings, optionally with external embeddings
780 9    $E_{\text{counts}} \leftarrow \text{Embed expression counts}$ 
781   Embed counts via MLP encoder, apply masking for task-specific
782   inputs
783 10   $E_{\text{cond}} \leftarrow \text{Embed conditions}$ 
784   Embed dataset, task, and perturbation information (gene or
785   chemical)
786 11  return  $E_{\text{gene}}, E_{\text{counts}}, E_{\text{cond}}$ 
787
788 12 Function EncodeConditions ( $E_{\text{cond}}$ ):
789 13  Apply bidirectional self-attention over condition tokens
790 14  Multiple layers of self-attention + FFN with RMSNorm
791 15  return  $H_{\text{cond}}$ 
792
793 16 Function CompressToLatents ( $E_{\text{input}}, H_{\text{cond}}$ ):
794 17  Initialize  $N_{\text{latent}}$  learnable latent tokens
795 18  Latent tokens:  $L_{\text{init}} \in \mathbb{R}^{N_{\text{latent}} \times d}$ 
796 19  Apply self-attention among latent tokens
797 20  Apply cross-attention: queries from latents, keys/values from  $E_{\text{input}}$ 
798 21  Apply cross-attention: queries from latents, keys/values from  $H_{\text{cond}}$ 
799 22  Decoder applies multiple layers with FFN and RMSNorm
800 23  return  $L$ 
801
802
803
804
805
806
807
808
809

```

---

**Algorithm 4: Task Head Application**


---

```

810
811
812 1 Function ApplyTaskHead( $L, \mathcal{T}$ ):
813   Input: Latent states  $L \in \mathbb{R}^{B \times N_{\text{latent}} \times d}$ , task  $\mathcal{T}$ 
814   Output: Task outputs  $\mathcal{O}$ , losses  $\mathcal{L}$ 
815 2    $H_{\text{genes}} \leftarrow \text{ProjectToGenes}(L)$ 
816   Project latent states to gene space via cross-attention
817   Output:  $H_{\text{genes}} \in \mathbb{R}^{G \times d}$ 
818 3   if  $\mathcal{T}$  is perturbation prediction then
819   |    $\mathcal{O} \leftarrow \text{ApplyPerturbationHead}(H_{\text{genes}})$ 
820   else
821   |   if  $\mathcal{T}$  is classification then
822   |   |    $\mathcal{O} \leftarrow \text{ApplyClassificationHead}(H_{\text{genes}})$ 
823   |   else
824   |   |
825   |   |    $\mathcal{L} \leftarrow \text{ComputeTaskLoss}(\mathcal{O}, \text{targets}, \mathcal{T})$ 
826   |   |   Compute task-specific loss (see Appendix J)
827   |   return  $\mathcal{O}, \mathcal{L}$ 
828
829 11 Function ProjectToGenes( $L$ ):
830   Apply projection cross-attention
831   Queries from gene embeddings, keys/values from latent states  $L$ 
832 12 return  $H_{\text{genes}}$ 
833
834 13 Function ApplyPerturbationHead( $H_{\text{genes}}$ ):
835   Apply cross-attention with gene token embeddings
836   Project to gene-level expression values
837 14 Output:  $\hat{X} \in \mathbb{R}^{B \times G}$ 
838   return  $\hat{X}$ 
839
840 15 Function ApplyClassificationHead( $H_{\text{genes}}$ ):
841   Aggregate  $H_{\text{genes}}$  into per-sample vector
842   Mean pooling or learned summarization:  $h \in \mathbb{R}^d$ 
843 16 Map  $h$  to class logits
844   Apply softmax to obtain class probabilities
845 17 Output:  $\hat{y} \in [0, 1]^{B \times C}$ 
846   return  $\hat{y}$ 
847
848 19 Function ComputeTaskLoss( $\mathcal{O}, \text{targets}, \mathcal{T}$ ):
849   if  $\mathcal{T}$  is perturbation prediction then
850   |   Compute distribution loss and bulk prediction loss
851   else
852   |   if  $\mathcal{T}$  is masked counts then
853   |   |   Compute MSE loss on masked and unmasked positions
854   |   else
855   |   |   Compute task-specific loss
856   |   return  $\mathcal{L}$ 
857 24

```

---

858  
859 The backbone algorithm processes batches of  $B$  cells, each with  $G$  genes, producing compressed  
860 latent representations of dimension  $N_{\text{latent}} \times d$ , where  $N_{\text{latent}} \ll G$  enables efficient processing of  
861 high-dimensional gene expression data. The decoder’s dual cross-attention mechanism allows it to  
862 selectively attend to relevant genes while conditioning on the experimental context. Task-specific  
863 processing and head application extend the backbone outputs to produce task-appropriate predictions  
and losses.

## B DATASETS AND VALIDATION SPLITS

This section provides detailed descriptions of the datasets used for training and evaluating the Funomics T0 model, including the specific validation split strategies employed for each task.

### B.1 PRETRAINING DATA

Datasets are split by cell line and listed in Table 4. These datasets are used for both the masked count task and the perturbation prediction task.

### B.2 VALIDATION STRATEGY

We evaluate genetic perturbation prediction on two diverse single-cell RNA-seq datasets:

- **HEPG2 (Nadig 2025):** This dataset (Nadig et al., 2025) consists of CRISPR-mediated perturbations in the HEPG2 cell line.
- **K562 (Replogle 2022):** A large-scale genome-wide CRISPRi screen in K562 cells (Replogle et al., 2022).

**Validation Strategy:** For these datasets, we employ two complementary splitting strategies to assess generalization:

1. **Leave-Perturbation-Out (LPO):** A subset of perturbations is held out during training. The model must predict the response to novel genetic interventions not seen in the training set.
2. **Leave-Cell-line-Out (LCO):** Where applicable, models are evaluated on cell lines entirely withheld from the training phase to test cross-context generalization.

Chemical perturbation performance is evaluated using the Tahoe datasets for HOP62 cell line. These datasets capture cellular responses to a variety of small-molecule compounds at multiple dosages.

**Validation Strategy:** Similar to genetic perturbations, we use a held-out compound strategy. A fixed percentage of chemical compounds (and all their associated dosages) are moved to the validation set, ensuring the model generalizes to unseen molecular structures.

### B.3 TISSUE ANNOTATION (GTEx v10)

For tissue annotation, we use the Gene Tissue Expression (GTEx) v10 dataset (Lonsdale et al., 2013), comprising 19,788 bulk RNA-seq samples from 946 donors across 54 tissue types.

**Validation Splits:** We define two specific validation splits to test different aspects of model robustness:

1. **Validation Purity Split (`val_purity`):** Designed to identify the most challenging samples where tissue-specific signals may be subtle or mixed. The split is constructed as follows:
  - **Preprocessing:** TPM-normalized counts are total-count normalized and log-transformed (`log1p`). The top 5,000 highly variable genes are selected.
  - **Dimensionality Reduction:** Data is standardized and projected onto 50 principal components.
  - **Purity Calculation:** For each sample, we compute the 10 nearest neighbors in PCA space. Purity is defined as the fraction of neighbors sharing the same tissue label.
  - **Assignment:** For each tissue (with  $\geq 100$  samples), we select 50 samples with the lowest purity (or up to 30% of the tissue’s total samples) for the validation set.
2. **Validation Tissue Split (`val_tissue`):** A tissue-balanced split that mimics standard evaluation protocols in prior literature (Yap et al., 2021). We select 50 samples per tissue type (up to 30% of available observations), excluding tissues with fewer than 100 samples.

Samples not included in either validation split are used for training (14,849 samples). Tissues with extremely low representation (Kidney-Medulla, Cervix-Ectocervix, Cervix-Endocervix, Fallopian Tube, Bladder) are excluded from validation to ensure stable metric calculation.

## C TASK-SPECIFIC DATA DETAILS

This section provides detailed information about data used for each specific task, including datasets, split strategies, and evaluation metrics.

Datasets are split by cell line. The following table lists all datasets used for perturbation prediction.

### C.1 ANNOTATION TASK DATA

**Tissue Annotation** In the tissue annotation task, the model predicts tissue labels from gene expression profiles. For this purpose, we employed the Gene Tissue Expression (GTEx) v10 dataset, composed of 19788 samples from 946 donors, spanning 54 distinct tissue types, profiled using the bulk RNA-sequencing technology. We used the Transcripts Per Million (TPM) normalized expression data instead of raw counts data to account for different gene lengths. As the number of samples per tissue is quite unbalanced, with a minimum of 11 observations per tissue, and maximum of 800, we excluded five tissues with less than 100 observations from the validation splits. The following tissues were excluded: Kidney - Medulla - 11, Cervix - Ectocervix - 24, Fallopian Tube - 29, Cervix - Endocervix - 23, Bladder - 77.

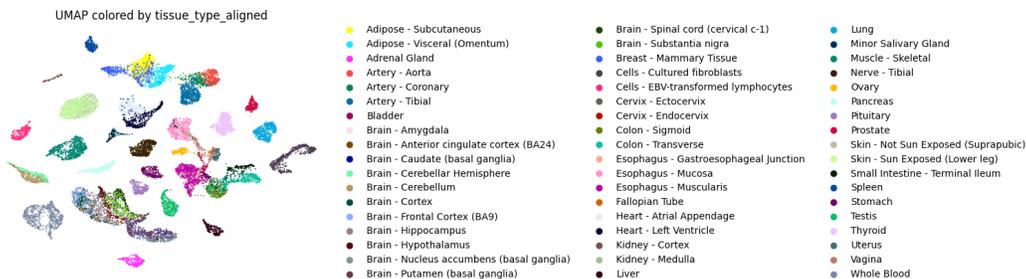


Figure 5: **UMAP visualization of GTEx v10 gene expression profiles colored by tissue type.** Each point represents a sample from the GTEx v10 bulk RNA-seq dataset, colored according to its annotated tissue label. This figure highlights the degree of separability and clustering structure among tissues in high-dimensional expression space.

For the tissue annotation task, we employ two complementary validation splits: `val_purity` (2426 samples) and `val_tissue` (2341 samples), with the remaining 14849 samples used for training.

Details on the GTEx tissue annotation dataset, including split definitions and procedures, are provided in subsection C.1.

## D ADDITIONAL RESULTS

Table 5 reports additional L1 metrics (Rank, Rank.T, and SMD) for genetic perturbation prediction.

Table 6 reports additional L1 metrics (Rank, Rank.T, and SMD) for chemical perturbation prediction on the Tahoe dataset.

## E BACKBONE ARCHITECTURE DETAILS

The Funomics T0 backbone implementation is based on T5Gemma (Zhang et al., 2025a) components for the encoder and decoder layers. T5Gemma (Zhang et al., 2025a) provides an efficient encoder-decoder architecture that is well-suited for processing structured inputs with cross-attention mechanisms. The architecture follows a set transformer design (Lee et al., 2019), treating inputs as

Table 4: Datasets used for perturbation prediction.

Dataset Name	Modality	Organism	Perturbation Type	Observations	Perturbations	Batches	Reference
<i>Genetic perturbation</i>							
VCC.H1_hESC_training_adata_unified_gene_set	scRNA-seq	<i>H. sapiens</i>	Genetic Knock Down	221,273	150	3	(Jaume et al.)
scPerturb_nadig_2025_JURKAT	scRNA-seq	<i>H. sapiens</i>	Genetic Knock Down	262,803	2,389	55	(Nadig et al., 2025)
scPerturb_nadig_2025_HEPG2	scRNA-seq	<i>H. sapiens</i>	Genetic Knock Down	145,427	2,389	56	(Nadig et al., 2025)
jiang_2025_AS49	scRNA-seq	<i>H. sapiens</i>	Genetic Knock Down	207,261	218	2	(Jiang et al., 2024)
jiang_2025_BXPC3	scRNA-seq	<i>H. sapiens</i>	Genetic Knock Down	314,758	218	2	(Jiang et al., 2024)
jiang_2025_HAP1	scRNA-seq	<i>H. sapiens</i>	Genetic Knock Down	277,261	218	2	(Jiang et al., 2024)
feng_2024_TargetedScreen_HPS10114i-eip1.1	scRNA-seq	<i>H. sapiens</i>	Genetic Knock Down	38,699	444	72	(Feng et al., 2024)
feng_2024_TargetedScreen_HPS10114i-zapf.3	scRNA-seq	<i>H. sapiens</i>	Genetic Knock Down	37,180	444	70	(Feng et al., 2024)
feng_2024_TargetedScreen_HPS10114i-kolf.2	scRNA-seq	<i>H. sapiens</i>	Genetic Knock Down	40,563	444	62	(Feng et al., 2024)
feng_2024_TargetedScreen_HPS10114i-kolf.3	scRNA-seq	<i>H. sapiens</i>	Genetic Knock Down	36,787	444	72	(Feng et al., 2024)
feng_2024_TargetedScreen_HPS10114i-ndw.4	scRNA-seq	<i>H. sapiens</i>	Genetic Knock Down	46,376	444	70	(Feng et al., 2024)
feng_2024_TargetedScreen_HPS1014i-oidk.2	scRNA-seq	<i>H. sapiens</i>	Genetic Knock Down	14,127	444	72	(Feng et al., 2024)
feng_2024_TargetedScreen_HPS1014i-oidk.5	scRNA-seq	<i>H. sapiens</i>	Genetic Knock Down	45,507	444	70	(Feng et al., 2024)
feng_2024_TargetedScreen_HPS1014i-riaj.1	scRNA-seq	<i>H. sapiens</i>	Genetic Knock Down	41,798	444	70	(Feng et al., 2024)
feng_2024_TargetedScreen_HPS1014i-riaj.3	scRNA-seq	<i>H. sapiens</i>	Genetic Knock Down	8,254	443	72	(Feng et al., 2024)
feng_2024_TargetedScreen_HPS1014i-paab.3	scRNA-seq	<i>H. sapiens</i>	Genetic Knock Down	28,918	444	70	(Feng et al., 2024)
feng_2024_TargetedScreen_HPS1014i-paab.4	scRNA-seq	<i>H. sapiens</i>	Genetic Knock Down	42,507	444	72	(Feng et al., 2024)
feng_2024_TargetedScreen_HPS1014i-ndw.1	scRNA-seq	<i>H. sapiens</i>	Genetic Knock Down	44,469	444	72	(Feng et al., 2024)
feng_2024_TargetedScreen_HPS10174i-ndw.4	scRNA-seq	<i>H. sapiens</i>	Genetic Knock Down	34,702	444	70	(Feng et al., 2024)
feng_2024_TargetedScreen_HPS10174i-pipw.4	scRNA-seq	<i>H. sapiens</i>	Genetic Knock Down	37,178	444	72	(Feng et al., 2024)
feng_2024_TargetedScreen_HPS10174i-pipw.5	scRNA-seq	<i>H. sapiens</i>	Genetic Knock Down	9,325	444	70	(Feng et al., 2024)
feng_2024_TargetedScreen_HPS101914i-efj.2	scRNA-seq	<i>H. sapiens</i>	Genetic Knock Down	47,558	444	72	(Feng et al., 2024)
feng_2024_TargetedScreen_HPS101914i-efj.3	scRNA-seq	<i>H. sapiens</i>	Genetic Knock Down	34,851	444	54	(Feng et al., 2024)
feng_2024_TargetedScreen_HPS11213i-toig.4	scRNA-seq	<i>H. sapiens</i>	Genetic Knock Down	16,732	444	72	(Feng et al., 2024)
feng_2024_TargetedScreen_HPS11213i-toig.6	scRNA-seq	<i>H. sapiens</i>	Genetic Knock Down	29,491	444	70	(Feng et al., 2024)
scplex_xe2_genetic_A172	scRNA-seq	<i>H. sapiens</i>	Genetic Knock Down	47,452	522	32	(Srivatsan et al., 2020)
scplex_xe2_genetic_T98G	scRNA-seq	<i>H. sapiens</i>	Genetic Knock Down	44,792	522	32	(Srivatsan et al., 2020)
scplex_xe2_genetic_LU87MG	scRNA-seq	<i>H. sapiens</i>	Genetic Knock Down	47,026	522	32	(Srivatsan et al., 2020)
scPerturb_Replogle2022_RPE1essential	scRNA-seq	<i>H. sapiens</i>	Genetic Knock Down	247,824	2,389	56	(Replogle et al., 2022)
scPerturb_Replogle2022_K562essential	scRNA-seq	<i>H. sapiens</i>	Genetic Knock Down	310,342	2,054	52	(Replogle et al., 2022)
scPerturb_Replogle2022_K562gwps	scRNA-seq	<i>H. sapiens</i>	Genetic Knock Down	1,988,657	9,854	267	(Replogle et al., 2022)
jiang_2025_HT29	scRNA-seq	<i>H. sapiens</i>	Genetic Knock Down	360,963	218	2	(Jiang et al., 2024)
jiang_2025_K562	scRNA-seq	<i>H. sapiens</i>	Genetic Knock Down	217,688	218	2	(Jiang et al., 2024)
jiang_2025_MCF7	scRNA-seq	<i>H. sapiens</i>	Genetic Knock Down	260,545	218	2	(Jiang et al., 2024)
xatlas_filtered_dual_guide_cells_HCT116	scRNA-seq	<i>H. sapiens</i>	Genetic Knock Down	349,927	2,000	109	(Moreno et al., 2021)
xatlas_filtered_dual_guide_cells_HCT116	scRNA-seq	<i>H. sapiens</i>	Genetic Knock Down	362,424	1,999	109	(Moreno et al., 2021)
xatlas_filtered_dual_guide_cells_HCT116	scRNA-seq	<i>H. sapiens</i>	Genetic Knock Down	344,738	1,998	109	(Moreno et al., 2021)
xatlas_filtered_dual_guide_cells_HCT116	scRNA-seq	<i>H. sapiens</i>	Genetic Knock Down	354,504	1,998	109	(Moreno et al., 2021)
xatlas_filtered_dual_guide_cells_HCT116	scRNA-seq	<i>H. sapiens</i>	Genetic Knock Down	349,785	2,000	109	(Moreno et al., 2021)
xatlas_filtered_dual_guide_cells_HCT116	scRNA-seq	<i>H. sapiens</i>	Genetic Knock Down	50,924	293	109	(Moreno et al., 2021)
xatlas_filtered_dual_guide_cells_HCT116	scRNA-seq	<i>H. sapiens</i>	Genetic Knock Down	352,573	2,000	109	(Moreno et al., 2021)
xatlas_filtered_dual_guide_cells_HCT116	scRNA-seq	<i>H. sapiens</i>	Genetic Knock Down	354,801	2,000	109	(Moreno et al., 2021)
xatlas_filtered_dual_guide_cells_HCT116	scRNA-seq	<i>H. sapiens</i>	Genetic Knock Down	357,281	2,000	109	(Moreno et al., 2021)
xatlas_filtered_dual_guide_cells_HCT116	scRNA-seq	<i>H. sapiens</i>	Genetic Knock Down	364,990	1,997	109	(Moreno et al., 2021)
xatlas_filtered_dual_guide_cells_HCT116	scRNA-seq	<i>H. sapiens</i>	Genetic Knock Down	165,777	0	109	(Moreno et al., 2021)
xatlas_filtered_dual_guide_cells_HEK293T	scRNA-seq	<i>H. sapiens</i>	Genetic Knock Down	233,951	1,000	223	(Moreno et al., 2021)
xatlas_filtered_dual_guide_cells_HEK293T	scRNA-seq	<i>H. sapiens</i>	Genetic Knock Down	239,735	1,000	223	(Moreno et al., 2021)
xatlas_filtered_dual_guide_cells_HEK293T	scRNA-seq	<i>H. sapiens</i>	Genetic Knock Down	240,981	999	223	(Moreno et al., 2021)
xatlas_filtered_dual_guide_cells_HEK293T	scRNA-seq	<i>H. sapiens</i>	Genetic Knock Down	218,838	0	223	(Moreno et al., 2021)
xatlas_filtered_dual_guide_cells_HEK293T	scRNA-seq	<i>H. sapiens</i>	Genetic Knock Down	234,823	1,000	223	(Moreno et al., 2021)
xatlas_filtered_dual_guide_cells_HEK293T	scRNA-seq	<i>H. sapiens</i>	Genetic Knock Down	232,650	1,000	223	(Moreno et al., 2021)
xatlas_filtered_dual_guide_cells_HEK293T	scRNA-seq	<i>H. sapiens</i>	Genetic Knock Down	249,132	1,000	223	(Moreno et al., 2021)
xatlas_filtered_dual_guide_cells_HEK293T	scRNA-seq	<i>H. sapiens</i>	Genetic Knock Down	232,787	999	223	(Moreno et al., 2021)
xatlas_filtered_dual_guide_cells_HEK293T	scRNA-seq	<i>H. sapiens</i>	Genetic Knock Down	237,874	998	223	(Moreno et al., 2021)
scPerturb_Replogle2022_K562gwps	scRNA-seq	<i>H. sapiens</i>	Genetic Knock Down	876,979	3,286	267	(Replogle et al., 2022)
xatlas_filtered_dual_guide_cells_HEK293T	scRNA-seq	<i>H. sapiens</i>	Genetic Knock Down	219,246	1,000	223	(Moreno et al., 2021)
xatlas_filtered_dual_guide_cells_HEK293T	scRNA-seq	<i>H. sapiens</i>	Genetic Knock Down	231,660	999	223	(Moreno et al., 2021)
scPerturb_Replogle2022_K562gwps	scRNA-seq	<i>H. sapiens</i>	Genetic Knock Down	668,691	3,283	267	(Replogle et al., 2022)
xatlas_filtered_dual_guide_cells_HEK293T	scRNA-seq	<i>H. sapiens</i>	Genetic Knock Down	237,919	999	223	(Moreno et al., 2021)
xatlas_filtered_dual_guide_cells_HEK293T	scRNA-seq	<i>H. sapiens</i>	Genetic Knock Down	234,524	1,000	223	(Moreno et al., 2021)
scPerturb_Replogle2022_K562gwps	scRNA-seq	<i>H. sapiens</i>	Genetic Knock Down	442,987	3,288	267	(Replogle et al., 2022)
xatlas_filtered_dual_guide_cells_HEK293T	scRNA-seq	<i>H. sapiens</i>	Genetic Knock Down	231,491	1,000	223	(Moreno et al., 2021)
xatlas_filtered_dual_guide_cells_HEK293T	scRNA-seq	<i>H. sapiens</i>	Genetic Knock Down	72,981	311	223	(Moreno et al., 2021)
xatlas_filtered_dual_guide_cells_HEK293T	scRNA-seq	<i>H. sapiens</i>	Genetic Knock Down	228,542	1,000	223	(Moreno et al., 2021)
xatlas_filtered_dual_guide_cells_HEK293T	scRNA-seq	<i>H. sapiens</i>	Genetic Knock Down	236,216	1,000	223	(Moreno et al., 2021)
xatlas_filtered_dual_guide_cells_HEK293T	scRNA-seq	<i>H. sapiens</i>	Genetic Knock Down	226,008	1,000	223	(Moreno et al., 2021)
xatlas_filtered_dual_guide_cells_HEK293T	scRNA-seq	<i>H. sapiens</i>	Genetic Knock Down	249,862	1,000	223	(Moreno et al., 2021)
xatlas_filtered_dual_guide_cells_HEK293T	scRNA-seq	<i>H. sapiens</i>	Genetic Knock Down	243,051	1,000	223	(Moreno et al., 2021)
PerturbAtlas_hESC_bulk_genetic_KD.h5ad	bulk RNA-seq	<i>H. sapiens</i>	Genetic Knock Down	18	1	2	(Zhang et al., 2025c)
<i>Chemical perturbation</i>							
scplex_xe2_chemical_A172	scRNA-seq	<i>H. sapiens</i>	Chemical	236,302	8	32	(Srivatsan et al., 2020)
scplex_xe2_chemical_T98G	scRNA-seq	<i>H. sapiens</i>	Chemical	212,789	8	32	(Srivatsan et al., 2020)
scplex_xe2_chemical_LU87MG	scRNA-seq	<i>H. sapiens</i>	Chemical	238,788	8	32	(Srivatsan et al., 2020)
scplex3_MCF7	scRNA-seq	<i>H. sapiens</i>	Chemical	344,837	752	16	(Srivatsan et al., 2020)
scplex3_K562	scRNA-seq	<i>H. sapiens</i>	Chemical	173,632	752	16	(Srivatsan et al., 2020)
scplex3_AS49	scRNA-seq	<i>H. sapiens</i>	Chemical	244,268	752	20	(Srivatsan et al., 2020)
taohe_100m_SW1271	scRNA-seq	<i>H. sapiens</i>	Chemical	209,283	1,131	42	(Zhang et al., 2025b)
taohe_100m_AS49	scRNA-seq	<i>H. sapiens</i>	Chemical	2,567,838	1,134	42	(Zhang et al., 2025b)
taohe_100m_SW1088	scRNA-seq	<i>H. sapiens</i>	Chemical	1,49,059	1,133	42	(Zhang et al., 2025b)
taohe_100m_AN3CA	scRNA-seq	<i>H. sapiens</i>	Chemical	566,078	1,134	42	(Zhang et al., 2025b)
taohe_100m_SW1417	scRNA-seq	<i>H. sapiens</i>	Chemical	1,921,742	1,134	42	(Zhang et al., 2025b)
taohe_100m_SKMEL2	scRNA-seq	<i>H. sapiens</i>	Chemical	1,548,711	1,133	42	(Zhang et al., 2025b)
taohe_100m_SW48	scRNA-seq	<i>H. sapiens</i>	Chemical	954,199	1,134	42	(Zhang et al., 2025b)
taohe_100m_SNU1	scRNA-seq	<i>H. sapiens</i>	Chemical	1,221,708	1,134	42	(Zhang et al., 2025b)
taohe_100m_A172	scRNA-seq	<i>H. sapiens</i>	Chemical	2,314,134	1,134	42	(Zhang et al., 2025b)
taohe_100m_ASP1	scRNA-seq	<i>H. sapiens</i>	Chemical	2,259,176	1,134	42	(Zhang et al., 2025b)
taohe_100m_BT474	scRNA-seq	<i>H. sapiens</i>	Chemical	1,271,918	1,133	42	(Zhang et al., 2025b)
taohe_100m_COLO205	scRNA-seq	<i>H. sapiens</i>	Chemical	1,456,805	1,134	42	(Zhang et al., 2025b)
taohe_100m_HCT15	scRNA-seq	<i>H. sapiens</i>	Chemical	1,500,121	1,134	42	(Zhang et al., 2025b)
taohe_100m_HEC1A	scRNA-seq	<i>H. sapiens</i>	Chemical	2,514,625	1,134	42	(Zhang et al., 2025b)
taohe_100m_HT29	scRNA-seq	<i>H. sapiens</i>	Chemical	2,171,036	1,134	42	(Zhang et al., 2025b)
taohe_100m_H5578T	scRNA-seq	<i>H. sapiens</i>	Chemical	1,725,286	1,134	42	(Zhang et al., 2025b)
taohe_100m_H5766T	scRNA-seq	<i>H. sapiens</i>	Chemical	2,535,736	1,133	42	(Zhang et al., 2025b)
taohe_100m_J82	scRNA-seq	<i>H. sapiens</i>	Chemical	2,342,982	1,134	42	(Zhang et al., 2025b)
taohe_100m_SNU423	scRNA-seq	<i>H. sapiens</i>	Chemical	1,501,700	1,134	42	(Zhang et al., 2025b)
taohe_100m_KATOIII	scRNA-seq	<i>H. sapiens</i>	Chemical	2,229,643	1,134	42	(Zhang et al., 2025b)
taohe_100m_L8180	scRNA-seq	<i>H. sapiens</i>	Chemical	1,828,299	1,134	42	(Zhang et al., 2025b)
taohe_100m_LOVO	scRNA-seq	<i>H. sapiens</i>	Chemical	3,013,246	1,134	42	(Zhang et al., 2025b)
taohe_100m_MIA-PACA2	scRNA-seq	<i>H. sapiens</i>	Chemical	2,432,647	1,133	42	(Zhang et al., 2025b)
taohe_100m_NCH460	scRNA-seq	<i>H. sapiens</i>	Chemical	5,736,238	1,134	42	(Zhang et al., 2025b)
taohe_100m_PAN1C1	scRNA-seq	<i>H. sapiens</i>	Chemical	4,089,886	1,133	42	(Zhang et al., 2025b)
taohe_100m_RKO	scRNA-seq	<i>H. sapiens</i>	Chemical	2,182,314	1,134	42	(Zhang et al., 2025b)
taohe_100m_SW480	scRNA-seq	<i>H. sapiens</i>	Chemical	6,040,371	1,134	42	(Zhang et al., 2025b)
taohe_100m_A427	scRNA-seq	<i>H. sapiens</i>	Chemical	1,430,114	1,134	42	(Zhang et al., 2025b)
taohe_100m_A498	scRNA-seq	<i>H. sapiens</i>	Chemical	2,737,714	1,134	42	(Zhang et al., 2025b)
taohe_100m_C33A	scRNA-seq	<i>H. sapiens</i>	Chemical	1,483,268	1,134	42	(Zhang et al., 2025b)
taohe_100m_C32	scRNA-seq	<i>H. sapiens</i>	Chemical	1,991,252	1,134	42	(Zhang et al., 2025b)
taohe_100m_C3A	scRNA-seq	<i>H. sapiens</i>	Chemical	1,244,192	1,134	42	(Zhang et al., 2025b)
taohe_100m_CFPAC1	scRNA-seq	<i>H. sapiens</i>	Chemical	2,404,632	1,134	42	(Zhang et al., 2025b)
taohe_100m_CHP121	scRNA-seq	<i>H. sapiens</i>	Chemical	537,791	1,134	42	(Zhang et al., 2025b)
taohe_100m_H4	scRNA-seq	<i>H. sapiens</i>	Chemical	1,042,066	1,133	42	(Zhang et al., 2025b)
taohe_100m_HOP62	scRNA-seq	<i>H. sapiens</i>	Chemical	2,307,302	1,134	42	(Zhang et al., 2025b)
taohe_100m_LOXIMV1	scRNA-seq	<i>H. sapiens</i>	Chemical	2,208,183	1,134	42	(Zhang et al., 2025b)
taohe_100m_NCH1573	scRNA-seq	<i>H. sapiens</i>	Chemical	2,016,991	1,134	42	(Zhang et al., 2025b)
taohe_100m_NCH1792	scRNA-seq	<i>H. sapiens</i>	Chemical	1,838			

1026 Table 5: Additional L1 metrics on genetic perturbation prediction (not shown in main results).

metric	HEPG2 (Nadig et al., 2025)				K562 (Replogle et al., 2022)			
	State	Funomics T0	mean pert	Bound*	State	Funomics T0	mean pert	Bound*
Rank (L1) ↓	0.22	0.09	1.00	0.02	0.23	0.21	1.00	0.00
Rank.T (L1) ↓	0.27	0.09	0.50	0.26	0.31	0.22	0.50	0.28
SMD (L1) ↓	50876.84	1834.25	14624.03	1675.30	46338.88	1419.17	10507.64	1362.63

1035 \*Bound is the better of technical duplicate and interpolated duplicate baselines.

1037 Table 6: Additional L1 metrics on chemical perturbation prediction (Tahoe).

metric	Tahoe (Zhang et al., 2025b)	
	State	Funomics
Rank (L1) ↓	0.00	0.10
Rank.T (L1) ↓	0.00	0.09
SMD (L1) ↓	11129.07	6392.14

## 1048 E.1 CONDITION ENCODER

1049 The condition encoder processes condition embeddings  $E_{\text{cond}} \in \mathbb{R}^{N_{\text{cond}} \times d}$  through multiple layers  
1050 of bidirectional self-attention, allowing each condition token to attend to all other condition tokens.  
1051 This enables the encoder to capture relationships between different condition components (e.g.,  
1052 perturbation target gene, dosage, task type) and produce encoded condition representations  $H_{\text{cond}} \in$   
1053  $\mathbb{R}^{N_{\text{cond}} \times d}$  that integrate information across all condition elements. The encoder uses standard  
1054 transformer layers with self-attention, feed-forward networks, RMSNorm normalization, and residual  
1055 connections, following the T5Gemma (Zhang et al., 2025a) architecture.

## 1058 E.2 DECODER

1059 The decoder operates on a fixed set of  $N_{\text{latent}}$  decoder input embeddings  $L_{\text{init}} \in \mathbb{R}^{N_{\text{latent}} \times d}$ , where  
1060  $N_{\text{latent}} \ll G$  (typically  $N_{\text{latent}} \approx 256$  while  $G \approx 19,000$ ). These decoder inputs are obtained by  
1061 embedding a fixed vocabulary of  $N_{\text{latent}}$  tokens through learned embedding lookups, rather than  
1062 randomly initialized parameters. The decoder uses a dual cross-attention mechanism to condition on  
1063 both gene expression data and experimental conditions.

1065 **Conditioning on Gene Expression:** The decoder applies cross-attention to the combined gene  
1066 embeddings  $E_{\text{input}} = E_{\text{gene}} + E_{\text{counts}}$ , where queries come from the latent tokens and keys/values  
1067 come from the gene embeddings. This allows the latent tokens to selectively attend to relevant genes  
1068 from the high-dimensional input, enabling the model to focus on genes that are most informative for  
1069 the prediction task.

1070 **Conditioning on Experimental Context:** The decoder also applies cross-attention to the encoded  
1071 condition embeddings  $H_{\text{cond}}$  from the encoder, enabling the model to condition its predictions on the  
1072 specific experimental context (e.g., which gene is perturbed, what dosage is used, what task is being  
1073 performed). This encoder-decoder cross-attention mechanism connects the condition encoder and  
1074 decoder, allowing condition information to guide the gene expression processing.

1075 **Information Flow:** Within each decoder layer, self-attention among latent tokens enables information  
1076 exchange, followed by the two cross-attention mechanisms (to genes and to conditions), and feed-  
1077 forward processing. The decoder produces compressed latent states  $L \in \mathbb{R}^{N_{\text{latent}} \times d}$  that integrate  
1078 information from both the gene expression input and the experimental conditions. These latent states  
1079 are then projected back to gene space via projection cross-attention (described in subsection E.3),  
producing gene-level hidden states  $H_{\text{genes}} \in \mathbb{R}^{G \times d}$  that serve as input to task-specific heads.

### 1080 E.3 PROJECTION CROSS-ATTENTION

1081  
1082 After the decoder produces compressed latent states  $L$ , a projection cross-attention mechanism maps  
1083 these latent representations back to gene space, producing gene-level hidden states  $H_{\text{genes}} \in \mathbb{R}^{G \times d}$ .  
1084 This projection step is essential for task-specific heads that operate on gene-level representations  
1085 rather than compressed latent tokens.

1086 The projection cross-attention uses gene embeddings  $E_{\text{gene}}$  as queries and the latent states  $L$  as keys  
1087 and values. Specifically, after applying RMSNorm normalization to the latent states, the projection  
1088 cross-attention computes:

$$1089 H_{\text{genes}} = \text{CrossAttention}(Q = E_{\text{gene}}, K = L', V = L'), \quad (1)$$

1091 where  $L' = \text{RMSNorm}(L)$  are the normalized latent states. This mechanism allows each gene  
1092 embedding to attend to relevant information from the compressed latent representation, effectively  
1093 decompressing the latent states back to the full gene space while preserving the information captured  
1094 during compression.

1095 The projection cross-attention enables task-specific heads to operate directly on gene-level repre-  
1096 sentations  $H_{\text{genes}}$ , which is particularly important for tasks like perturbation prediction that require  
1097 gene-wise predictions. This design maintains the computational efficiency of the compressed latent  
1098 representation during processing while providing full-gene expressivity for downstream task heads.  
1099

### 1100 E.4 NORMALIZATION

1101 Gene token embeddings, counts embeddings, and all intermediate representations are normalized  
1102 using RMSNorm (Root Mean Square Layer Normalization). For an embedding  $E \in \mathbb{R}^d$ , RMSNorm  
1103 is defined as:

$$1104 E' = \text{RMSNorm}(E) = \frac{E}{\sqrt{\frac{1}{d} \sum_{i=1}^d E_i^2 + \epsilon}} \odot \gamma, \quad (2)$$

1105  
1106 where  $\gamma \in \mathbb{R}^d$  is a learnable scale parameter and  $\epsilon$  is a small constant for numerical stability.  
1107 RMSNorm is computationally efficient and has been shown to work well in transformer architectures.  
1108 Detailed descriptions of the embedding modules (gene token embeddings, counts embeddings, and  
1109 condition embeddings) are provided in Appendix H.  
1110  
1111

### 1112 E.5 ARCHITECTURAL CHOICES

1113  
1114 **Set Transformer Design:** The architecture treats inputs as unordered sets without positional embed-  
1115 dings, making it permutation-invariant. This is crucial for gene expression data where gene order is  
1116 arbitrary and should not affect predictions.

1117 **Dimensionality Reduction:** By choosing  $N_{\text{latent}} \ll G$ , the model achieves significant dimension-  
1118 ality reduction, making it computationally tractable while preserving essential information. The  
1119 compressed latent representation enables efficient processing of high-dimensional gene expression  
1120 profiles.  
1121

1122 **T5Gemma Components:** The use of T5Gemma components provides proven architectural patterns  
1123 for encoder-decoder models, including efficient attention mechanisms and normalization strategies  
1124 that have been validated in large-scale language models.  
1125

## 1126 F TRAINING DETAILS

1127  
1128 This section documents the concrete training configuration used for our reported model checkpoints,  
1129 focusing on details needed for reproducibility. Dataset and task definitions are described elsewhere in  
1130 the appendix.  
1131

### 1132 F.1 MODEL CONFIGURATION (CHOSEN HYPERPARAMETERS)

1133 We use the following model hyperparameters:

- **Number of latent tokens:**  $N_{\text{latent}} = 512$ .
- **Hidden dimension of latents in the decoder:** 1024.
- **Attention heads:** 8.
- **Attention dimension:** 128.
- **Decoder depth:** 16 layers.
- **Dropout:** 0.05.
- **Attention dropout:** 0.05.

## F.2 PRETRAINING RUN CONFIGURATION

Pretraining was run with the following settings.

- **Objectives:** pretraining jointly mixes count prediction and perturbation prediction tasks.
- **Batch size:** 64.
- **Learning rate:**  $1 \times 10^{-4}$  with a constant schedule and warmup fraction 0.05.
- **Training length:** 1M optimization steps.
- **Compute:** 7 days on 16 NVIDIA H100 GPUs (approximately 3000 GPU hours).
- **Implementation:** PyTorch with FlashAttention-2 enabled.
- **Pointers:** task definitions and losses are in Appendix I and Appendix J; datasets and split strategies are in Appendix B.

## G WEIGHT COMPUTATION FOR WMSE

Following the calibrated metrics framework Mejia et al. (2025) weights  $w_i$  are used in WMSE to prioritize genes of biological importance. However, we identified a problem with the standard weight computation approach from Mejia et al. (2025) To address this limitation, we propose stratified weights that provide improved gene prioritization.

## H EMBEDDINGS

This section provides detailed descriptions of all embedding modules used in the Funomics T0 model, including gene token embeddings, counts embeddings, and condition embeddings for experimental contexts.

### H.1 GENE TOKEN EMBEDDINGS

Each gene is represented by a learnable token embedding  $E_{\text{gene}} \in \mathbb{R}^{G \times d}$ , where  $G$  is the number of genes and  $d$  is the embedding dimension. Gene token embeddings are randomly initialized and learned during training. The embedding matrix  $E_{\text{gene}}$  provides then a representation of each gene in the unified gene set, enabling the model to learn gene-specific features.

Optionally, external gene embeddings can be integrated to provide prior knowledge. When external embeddings  $E_{\text{external}} \in \mathbb{R}^{G \times d_{\text{ext}}}$  are available, they are projected through a linear layer to match the model’s embedding dimension:  $E_{\text{projected}} = E_{\text{external}} W_{\text{proj}} + b_{\text{proj}}$ , where  $W_{\text{proj}} \in \mathbb{R}^{d_{\text{ext}} \times d}$  and  $b_{\text{proj}} \in \mathbb{R}^d$  are learnable parameters. The projected embeddings are then added to the token embeddings:  $E_{\text{gene}} = E_{\text{gene}} + E_{\text{projected}}$ , similarly to the positional embeddings used in classical transformer architecture (Vaswani et al., 2017).

### H.2 COUNTS EMBEDDINGS

**Binning Approach:** Following Adduri et al. (2025), continuous counts  $x_{b,g} \in \mathbb{R}$  are embedded using a learned binning function that maps count values to a weighted combination of bin embeddings. The binning function uses two linear layers with a LeakyReLU activation to produce soft weights over

bins, which are then used to compute a weighted sum of learnable bin embeddings. This approach provides a discrete-like representation that is well-suited for attention mechanisms and can capture non-linear relationships through the learned bin boundaries. Special mask and pad tokens are used for masked and padded positions respectively.

### H.3 CONDITION EMBEDDINGS

Condition embeddings encode experimental conditions and task information, allowing the model to condition its predictions on the specific context. All condition embeddings are concatenated into a condition sequence  $E_{\text{cond}} \in \mathbb{R}^{N_{\text{cond}} \times d}$ , where  $N_{\text{cond}}$  varies based on the perturbation type and task.

#### H.3.1 GENETIC PERTURBATION EMBEDDINGS

For genetic perturbations, the perturbation target gene is embedded using the same gene token embedding procedure as the input gene embeddings. Specifically, if the perturbation targets gene  $g_{\text{target}}$ , the condition embedding is  $E_{\text{genetic}} = E_{\text{gene}}[g_{\text{target}}] \in \mathbb{R}^d$ , where  $E_{\text{gene}}$  is the same gene token embedding matrix used for input genes. This ensures consistent representation of genes across the model, enabling the model to leverage learned gene relationships when predicting perturbation effects.

#### H.3.2 CHEMICAL PERTURBATION EMBEDDINGS

For chemical perturbations, multiple components are combined to form the condition embedding. Molecular embeddings  $E_{\text{mol}} \in \mathbb{R}^{d_{\text{mol}}}$  are obtained from external sources (e.g., molecular graph encoders or pre-trained chemical embeddings) and projected through an MLP to match the embedding dimension:  $E_{\text{mol,proj}} = \text{MLP}_{\text{mol}}(E_{\text{mol}}) \in \mathbb{R}^d$ . Dosage information is embedded through a learnable dosage embedding:  $E_{\text{dosage}} = E_{\text{dosage_emb}}[\text{bin}(d)] \in \mathbb{R}^d$ , where  $d$  is the dosage value and bin maps it to a discrete dosage bin. When available, molecules’s target gene embeddings  $E_{\text{target}} \in \mathbb{R}^d$  are included using the gene token embedding procedure. The final chemical perturbation embedding is constructed as:  $E_{\text{chemical}} = [E_{\text{mol,proj}}; E_{\text{dosage}}; E_{\text{target}}] \in \mathbb{R}^{N_{\text{chem}} \times d}$ , where  $N_{\text{chem}}$  depends on whether gene target information for a drug is available.

#### H.3.3 TASK EMBEDDINGS

Task embeddings are learnable embeddings that encode the task type, allowing the model to condition its behavior on the specific prediction task. Given a vocabulary of  $N_{\text{tasks}}$  task types (e.g., perturbation prediction, unmasking, annotation), each task  $t$  is represented by a learnable embedding  $E_{\text{task}}[t] \in \mathbb{R}^d$ . The task embedding vocabulary size  $N_{\text{tasks}}$  and embedding dimension  $d$  are hyperparameters that determine the model’s capacity to distinguish between different task types.

## I TASK DESCRIPTIONS

This section provides detailed descriptions of each task implemented in the Funomics T0 model, including task formulations, loss functions, and references to task-specific data details.

### I.1 PERTURBATION PREDICTIONS

#### I.1.1 GENETIC PERTURBATIONS

In the genetic perturbation prediction task, the model predicts perturbed cell expression profiles given control cell counts and genetic perturbation conditions. For a batch of  $B$  cell groups, each containing  $N$  cells, the model receives control counts  $X_{\text{ctrl}} \in \mathbb{R}^{B \times G}$  and produces predictions  $\hat{X} \in \mathbb{R}^{B \times G}$  for the perturbed state, where  $G$  is the number of genes. The target perturbed counts are  $X_{\text{pert}} \in \mathbb{R}^{B \times G}$ . Predictions are optionally normalized with scanpy’s `log1p(normalize_total)` and can be computed in residual mode by adding control means or average perturbation effects.

Task-specific data details, datasets, and split strategies are provided in subsection B.2.

**Loss Function** For genetic perturbation prediction, we use a combination of Maximum Mean Discrepancy (MMD) loss and bulk mean squared error (MSE) loss. The MMD term  $\mathcal{L}_{\text{MMD}} = \text{MMD}^2(\hat{X}, X_{\text{pert}})$  compares the distributions of predicted and target perturbed expressions, computed using the `geomloss` package with an energy kernel Bunne et al. (2023). Crucially, during training, each batch is constructed to contain all cells from the same cell line and perturbation condition, allowing the MMD loss to be computed over the entire set of predicted and observed cells for that condition. This distributional approach is necessary because there is no cell-level correspondence between control and perturbed populations due to the process of gene expression value measurement which is destructive for the cells. Therefore, rather than using point-wise losses that would require such matching, we measure distributional similarity between the predicted and observed perturbed cell sets. This loss captures the full distributional similarity between predicted and observed cell populations, preserving heterogeneity in perturbation responses. The bulk MSE loss  $\mathcal{L}_{\text{bulk}} = \|\hat{X} - \bar{X}_{\text{pert}}\|^2$  compares mean expressions across cells, where  $\hat{X} = \frac{1}{N} \sum_{n=1}^N \hat{X}_{:,n}$  and  $\bar{X}_{\text{pert}} = \frac{1}{N} \sum_{n=1}^N X_{\text{pert},:,n}$ , ensuring accurate prediction of average perturbation effects. The total perturbation loss is  $\mathcal{L}_{\text{pert}} = 0.2 \cdot \mathcal{L}_{\text{MMD}} + 10 \cdot \mathcal{L}_{\text{bulk}}$ . Optionally, gene-specific weights can be applied to emphasize differentially expressed genes, with weight computation details provided in Appendix G. A detailed explanation of the MMD formulation is provided in subsection J.1, and additional loss function details are described in Appendix J.

### I.1.2 CHEMICAL PERTURBATIONS

In the chemical perturbation prediction task, the model predicts perturbed cell expression profiles given control cell counts and chemical perturbation conditions (including compound identity and dosage). The formulation is similar to genetic perturbations, but condition embeddings incorporate molecular representations and dosage information.

**Loss Function** For chemical perturbation prediction, we use the same loss function as for genetic perturbations: a combination of Maximum Mean Discrepancy (MMD) loss and bulk mean squared error (MSE) loss. As with genetic perturbations, batches are constructed to contain all cells from the same cell line and perturbation condition, enabling distributional comparison via MMD (see subsection I.1.1 for the rationale). The MMD term  $\mathcal{L}_{\text{MMD}} = \text{MMD}^2(\hat{X}, X_{\text{pert}})$  compares the distributions of predicted and target perturbed expressions, computed using the `geomloss` package with an energy kernel Bunne et al. (2023). The bulk MSE loss  $\mathcal{L}_{\text{bulk}} = \|\hat{X} - \bar{X}_{\text{pert}}\|^2$  compares mean expressions across cells. The total perturbation loss is  $\mathcal{L}_{\text{pert}} = 0.2 \cdot \mathcal{L}_{\text{MMD}} + 10 \cdot \mathcal{L}_{\text{bulk}}$ . Optionally, gene-specific weights can be applied to emphasize differentially expressed genes, with weight computation details provided in Appendix G. A detailed explanation of the MMD formulation is provided in subsection J.1, and additional loss function details are described in Appendix J.

### I.2 UNMASKING PERTURBATION PREDICTION

In the unmasking perturbation prediction task, a subset of genes is masked during training, and the model predicts expression values for the masked genes. The model receives control counts  $X_{\text{ctrl}} \in \mathbb{R}^{B \times G}$  with a randomly masked subset of genes, and produces predictions  $\hat{X} \in \mathbb{R}^{B \times G}$  for all genes, where  $B$  is the batch size and  $G$  is the number of genes.

We perform masking directly on the input count matrix  $X_{\text{ctrl}} \in \mathbb{R}^{B \times G}$ : a binary mask  $M \in \{0, 1\}^{B \times G}$  is defined where  $M_{b,g} = 1$  indicates that the count value for gene  $g$  in batch  $b$  is masked (not observed), and  $M_{b,g} = 0$  indicates it is observed. For each masked entry in  $X_{\text{ctrl}}$ , the corresponding value is replaced during count embedding with a special mask token embedding. The model is then trained to predict the original observed count value at every masked position.

We employ a random masking strategy. For random masking, each element  $M_{b,g}$  is drawn from a Bernoulli distribution with parameter  $p_{\text{mask}}$ , i.e.,  $M_{b,g} \sim \text{Bernoulli}(p_{\text{mask}})$  for all  $b, g$ . The masking probability is task-specific and applied uniformly across genes.

**Loss Function** For unmasking tasks, we use MSE loss computed on masked positions during training, where the model is required to predict the original expression values. Given a binary mask  $M \in \{0, 1\}^{B \times G}$  where  $M_{b,g} = 1$  indicates that gene  $g$  in batch  $b$  is masked, the unmasking loss

is:  $\mathcal{L}_{\text{mask}} = \frac{1}{|\mathcal{M}|} \sum_{(b,g) \in \mathcal{M}} (\hat{X}_{b,g} - X_{\text{ctrl},b,g})^2$ , where  $\mathcal{M} = \{(b, g) : M_{b,g} = 1\}$  is the set of masked positions. This formulation encourages the model to learn to reconstruct expression values from partial observations, improving its ability to handle missing data and learn robust gene-gene relationships. Additional details are provided in Appendix J.

### I.3 ANNOTATION TASKS

#### I.3.1 TISSUE ANNOTATION

In the tissue annotation task, the goal is to predict the tissue of origin of a given sample, based on its expression profile. The formulation is similar to cell type annotation, but with tissue-specific labels.

Task-specific data details, datasets, and split strategies are provided in subsection B.3.

**Loss Function** For annotation tasks (cell type and tissue classification), we use cross-entropy loss to train the model for multi-class classification. The loss is defined as:  $\mathcal{L}_{\text{annot}} = -\frac{1}{B} \sum_{b=1}^B \sum_{c=1}^C y_{b,c} \log(\hat{y}_{b,c})$ , where  $y_{b,c} \in \{0, 1\}$  is the true one-hot label indicating whether cell  $b$  belongs to class  $c$ ,  $\hat{y}_{b,c} \in [0, 1]$  is the predicted probability for class  $c$  (obtained via softmax),  $B$  is the batch size, and  $C$  is the number of classes. This loss encourages the model to assign high probability to the correct class while minimizing probability for incorrect classes. Additional details are provided in Appendix J.

## J LOSS FUNCTION DETAILS

### J.1 MAXIMUM MEAN DISCREPANCY

Given samples  $\{x_i\}_{i=1}^m \sim P$  and  $\{y_j\}_{j=1}^n \sim Q$ , the squared Maximum Mean Discrepancy (MMD) between the two distributions is estimated empirically as

$$\text{MMD}^2(P, Q) = \frac{1}{m^2} \sum_{i,j=1}^m k(x_i, x_j) + \frac{1}{n^2} \sum_{i,j=1}^n k(y_i, y_j) - \frac{2}{mn} \sum_{i=1}^m \sum_{j=1}^n k(x_i, y_j), \quad (3)$$

where  $k(\cdot, \cdot)$  denotes a kernel function that determines the geometry of the comparison.

In our implementation, we adopt the *energy kernel*

$$k(x, y) = -\|x - y\|, \quad (4)$$

which is *conditionally negative definite* (CND). Although this kernel is not positive definite and therefore does not correspond to an inner product in a reproducing kernel Hilbert space (RKHS), it induces a well-defined metric between probability distributions. In particular, the resulting MMD coincides exactly with the *Energy Distance*, a classical statistical metric defined by

$$D_{\text{Energy}}(P, Q) = 2 \mathbb{E}\|X - Y\| - \mathbb{E}\|X - X'\| - \mathbb{E}\|Y - Y'\|. \quad (5)$$

Using a CND kernel ensures that the empirical estimate above defines a proper probability metric: it is non-negative, equals zero if and only if  $P = Q$ , and satisfies the triangle inequality. Compared to positive-definite kernels such as the RBF, the energy kernel requires no bandwidth hyperparameters and captures global geometric discrepancies through pairwise Euclidean distances.

### J.2 BULK MSE LOSS

The bulk MSE loss compares mean expressions across cells:

$$\mathcal{L}_{\text{bulk}} = \|\bar{\hat{X}} - \bar{X}_{\text{pert}}\|^2, \quad (6)$$

where  $\bar{\hat{X}} = \frac{1}{N} \sum_{n=1}^N \hat{X}_{:,n}$  and  $\bar{X}_{\text{pert}} = \frac{1}{N} \sum_{n=1}^N X_{\text{pert},:,n}$ .

### 1350 J.3 WEIGHTED MSE LOSS

1351  
1352 The weighted MSE loss corresponds to the Weighted Mean Squared Error (WMSE) metric defined  
1353 in subsection K.5. When used as a training objective, weighted MSE can be leveraged to drive the  
1354 model towards focusing more on accurate predictions for genes that are significantly affected by  
1355 perturbations. By assigning higher weights to genes with substantial perturbation-induced changes,  
1356 the loss function emphasizes learning to predict these biologically important responses while still  
1357 maintaining supervision across the full transcriptome.

### 1358 J.4 UNMASKING TASK LOSS

1359  
1360 For unmasking tasks, the model predicts expression values for masked genes. Given a binary mask  
1361  $M \in \{0, 1\}^{B \times G}$  where  $M_{b,g} = 1$  indicates that gene  $g$  in batch  $b$  is masked, the unmasking loss is  
1362 computed only on masked positions:

$$1363 \mathcal{L}_{\text{mask}} = \frac{1}{|\mathcal{M}|} \sum_{(b,g) \in \mathcal{M}} (\hat{X}_{b,g} - X_{\text{ctrl},b,g})^2, \quad (7)$$

1364  
1365 where  $\mathcal{M} = \{(b, g) : M_{b,g} = 1\}$  is the set of masked positions,  $\hat{X}_{b,g}$  is the predicted expression value,  
1366 and  $X_{\text{ctrl},b,g}$  is the ground truth expression value. This formulation trains the model to reconstruct  
1367 missing expression values from observed genes, improving its ability to learn gene-gene relationships  
1368 and handle incomplete data.  
1369  
1370

### 1371 J.5 ANNOTATION TASK LOSS

1372  
1373 For annotation tasks (cell type and tissue classification), we use cross-entropy loss. Given true  
1374 one-hot labels  $y \in \{0, 1\}^{B \times C}$  and predicted probabilities  $\hat{y} \in [0, 1]^{B \times C}$  (obtained via softmax), the  
1375 annotation loss is:

$$1376 \mathcal{L}_{\text{annot}} = -\frac{1}{B} \sum_{b=1}^B \sum_{c=1}^C y_{b,c} \log(\hat{y}_{b,c}), \quad (8)$$

1377  
1378 where  $B$  is the batch size and  $C$  is the number of classes. The predicted probabilities are computed as  
1379

$$1380 \hat{y}_{b,c} = \frac{\exp(z_{b,c})}{\sum_{c'=1}^C \exp(z_{b,c'})},$$

1381  
1382 where  $z_{b,c}$  are the logits from the classification head.  
1383

### 1384 J.6 MULTI-TASK LOSS WEIGHTING

1385  
1386 When training on multiple tasks simultaneously, the total loss is a weighted combination of task-  
1387 specific losses. Given tasks  $t \in \{1, 2, \dots, T\}$  with corresponding losses  $\mathcal{L}_t$  and weights  $\lambda_t$ , the total  
1388 loss is:

$$1389 \mathcal{L}_{\text{total}} = \sum_{t=1}^T \lambda_t \mathcal{L}_t. \quad (9)$$

1390  
1391 The weights  $\lambda_t$  can be set manually based on task importance, learned adaptively during training, or  
1392 determined through hyperparameter search. For perturbation prediction tasks, gene-specific weights  
1393 can be applied within the loss terms.  
1394

## 1395 K METRICS

### 1396 K.1 PREPROCESSING OF PREDICTIONS AND TARGETS

1397  
1398 Before computing metrics, predictions and targets may be preprocessed using various transforma-  
1399 tions to improve metric sensitivity and interpretability. We establish the following notation for  
1400 preprocessing operations:  
1401

1402 **Base (no preprocessing):** Raw expression values are used without transformation, denoted as “–” in  
1403 tables.

1404 **Control centering** ( $\Delta_{\text{ctrl}}$ ): Expression values are centered by subtracting the mean control expres-  
 1405 sion:

$$1406 \mathbf{x}_{\Delta_{\text{ctrl}}} = \mathbf{x} - \bar{\mathbf{x}}_{\text{ctrl}}, \quad (10)$$

1407 where  $\bar{\mathbf{x}}_{\text{ctrl}}$  is the mean expression across all control cells. This transformation removes baseline  
 1408 expression levels and emphasizes perturbation-specific changes.

1409 **Train perturbation centering** ( $\Delta_{\text{train}}$ ): Expression values are centered by subtracting the mean  
 1410 training perturbation expression:

$$1411 \mathbf{x}_{\Delta_{\text{train}}} = \mathbf{x} - \bar{\mathbf{x}}_{\text{train}}, \quad (11)$$

1412 where  $\bar{\mathbf{x}}_{\text{train}}$  is the mean expression across all training perturbations. This approach increases the  
 1413 dynamic range of metrics, particularly for datasets involving essential genes where perturbations may  
 1414 share effects on basic cellular processes Wu et al. (2025).  
 1415

## 1416 K.2 COSINE SIMILARITY

1418 Cosine similarity measures the angular similarity between two vectors, providing a scale-invariant  
 1419 measure of directional agreement. We use two variants of cosine similarity depending on the  
 1420 evaluation context.

### 1421 K.2.1 COSINE SIMILARITY

1422 Regular cosine similarity is computed between predicted and ground truth perturbation response  
 1423 vectors without any filtering. For a predicted vector  $\hat{\mathbf{x}}^i$  and ground truth vector  $\mathbf{x}^i$  for perturbation  $i$ ,  
 1424 cosine similarity is defined as:

$$1425 \text{cosine}(\hat{\mathbf{x}}^i, \mathbf{x}^i) = \frac{\hat{\mathbf{x}}^i \cdot \mathbf{x}^i}{\|\hat{\mathbf{x}}^i\|_2 \|\mathbf{x}^i\|_2}, \quad (12)$$

1426 where  $\cdot$  denotes the dot product and  $\|\cdot\|_2$  is the L2 norm. This metric ranges from  $-1$  to  $1$ , with  
 1427 values closer to  $1$  indicating better directional agreement. Higher values indicate better performance.  
 1428

### 1429 K.2.2 COSINE SIMILARITY FOR RESPONSE DIRECTION

1430 Cosine similarity for response direction evaluates the directional agreement specifically for pertur-  
 1431 bations with significant effect sizes. This variant focuses on perturbations that produce meaningful  
 1432 biological changes, filtering out perturbations with negligible effects.

1433 When calculating cosine similarity for response direction, a specific centering step is employed  
 1434 to improve the metric’s effectiveness (Littman et al., 2025). The average log-fold change (logFC)  
 1435 perturbation response in the training dataset is subtracted from both the prediction and the ground  
 1436 truth:

$$1437 \text{cosine}_{\text{response}}(\hat{\mathbf{x}}^i, \mathbf{x}^i) = \frac{(\hat{\mathbf{x}}^i - \bar{\mathbf{x}}_{\text{train}}) \cdot (\mathbf{x}^i - \bar{\mathbf{x}}_{\text{train}})}{\|(\hat{\mathbf{x}}^i - \bar{\mathbf{x}}_{\text{train}})\|_2 \|(\mathbf{x}^i - \bar{\mathbf{x}}_{\text{train}})\|_2}, \quad (13)$$

1438 where  $\bar{\mathbf{x}}_{\text{train}}$  is the mean perturbation response across the training dataset. This approach of centering  
 1439 at the mean perturbation response is adopted to increase the dynamic range of the metric (Littman  
 1440 et al., 2025). This centering is particularly important for datasets involving essential genes, where  
 1441 perturbations might share effects on basic cellular processes. This centering strategy deviates from  
 1442 conventional approaches used in prior work (Littman et al., 2025).  
 1443

## 1444 K.3 RANK METRICS

1445 Rank-based metrics measure perturbation discrimination by evaluating how model predictions com-  
 1446 pare to observations, following (Wu et al., 2025). This approach avoids the limitations of existing  
 1447 information retrieval metrics such as mean reciprocal rank, which would require a specific desired  
 1448 cell state to create rankings, potentially biasing the evaluation. Rank metrics are computed on a  
 1449 per-perturbation basis and use a generic distance metric  $\text{dist}(\cdot, \cdot)$  that can be L1 (Manhattan distance),  
 1450 L2 (Euclidean distance), or cosine similarity. The choice of distance metric can significantly affect  
 1451 the ranking (Liu et al., 2025). All rank metrics yield values between  $0$  and  $1$ , where  $0$  is a perfect  
 1452 score and  $0.5$  is the expected score of a random prediction. Lower values indicate better performance,  
 1453 with a value of  $0$  indicating perfect discrimination.  
 1454

### 1458 K.3.1 RANK<sub>average</sub>

1459  
1460 The Rank<sub>average</sub> metric measures perturbation discrimination by evaluating, for a given observed  
1461 perturbation, how close the model prediction is to the observation compared to predictions made for  
1462 other perturbations. For a set of  $p$  perturbations, Rank<sub>average</sub> is defined as:

$$1463 \text{Rank}_{\text{average}} := \frac{1}{p} \sum_{i=1}^p \text{rank}(\hat{x}^i), \quad (14)$$

1464 where  $\hat{x}^i$  is the predicted response for perturbation  $i$ , and the rank of prediction  $i$  is computed as:

$$1465 \text{rank}(\hat{x}^i) := \frac{1}{p-1} \sum_{j=1, j \neq i}^p \mathbb{I}(\text{dist}(\hat{x}^j, x^i) \leq \text{dist}(\hat{x}^i, x^i)), \quad (15)$$

1466 where the sum is over all  $j$  from 1 to  $p$  such that  $j \neq i$ ,  $x^i$  is the true (average) expression value  
1467 for perturbation  $i$ , and  $\mathbb{I}(\cdot)$  is the indicator function. The rank of prediction  $i$  represents the fraction  
1468 of other predictions that are at least as close to the true response  $x^i$  as the prediction  $\hat{x}^i$  itself.  
1469 Rank<sub>average</sub> is robust to scaling of predictions: adding a constant factor to all model predictions will  
1470 not affect the metric, as it only depends on the relative ordering of predictions with respect to each  
1471 true response. This makes Rank<sub>average</sub> particularly useful for assessing whether model predictions  
1472 are ordered correctly relative to the ground truth perturbations. The calibration of this rank metric  
1473 was investigated by Wu et al. (2025).  
1474  
1475  
1476  
1477

### 1478 K.3.2 RANK.T<sub>average</sub>

1479 The Rank.T<sub>average</sub> metric is a transposed variant of Rank<sub>average</sub> that measures perturbation discrimi-  
1480 nation by evaluating, for a given predicted perturbation response, how close it is to its corresponding  
1481 observation compared to other observations. For a set of  $p$  perturbations, Rank.T<sub>average</sub> is defined as:

$$1482 \text{Rank.T}_{\text{average}} := \frac{1}{p} \sum_{i=1}^p \text{rankT}(\hat{x}^i), \quad (16)$$

1483 where  $\hat{x}^i$  is the predicted response for perturbation  $i$ , and the rankT of prediction  $i$  is computed as:

$$1484 \text{rankT}(\hat{x}^i) := \frac{1}{p-1} \sum_{j=1, j \neq i}^p \mathbb{I}(\text{dist}(\hat{x}^i, x^j) \leq \text{dist}(\hat{x}^i, x^i)), \quad (17)$$

1485 where the sum is over all  $j$  from 1 to  $p$  such that  $j \neq i$ ,  $x^i$  is the true (average) expression value for  
1486 perturbation  $i$ , and  $\mathbb{I}(\cdot)$  is the indicator function. Unlike Rank<sub>average</sub> K.3.1, which compares different  
1487 predictions to the same true response, Rank.T<sub>average</sub> compares a single prediction to different true  
1488 responses. The rankT of prediction  $i$  represents the fraction of other true responses that are at least as  
1489 close to the prediction  $\hat{x}^i$  as the true response  $x^i$  itself. Unlike Rank<sub>average</sub>, Rank.T<sub>average</sub> is sensitive  
1490 to the scaling of predictions: adding a constant factor to all model predictions will affect the metric,  
1491 as it directly compares the absolute distances between predictions and true responses. This makes  
1492 Rank.T<sub>average</sub> useful for users who want a stricter assessment of whether the predictions are similar  
1493 to the ground truth perturbations, beyond just the ordering of the perturbations.  
1494  
1495  
1496  
1497  
1498

## 1499 K.4 SIMILARITY MATRIX DISTANCE

1500 The Similarity Matrix Distance (SMD) metric measures the preservation of relations between per-  
1501 turbations by quantifying the discrepancy between predicted and ground truth similarity matrices,  
1502 computed on pseudobulked data, following Wu et al. (2025). A similarity matrix  $S \in \mathbb{R}^{n \times n}$  cap-  
1503 tures pairwise distances for  $n$  perturbations, where each entry  $s_{ij}$  represents the distance between  
1504 perturbation  $i$  and perturbation  $j$ , computed from their aggregated expression vectors  $x^{(i)}, x^{(j)}$  using  
1505 a generic distance metric  $\text{dist}(\cdot, \cdot)$  that can be L1 (Manhattan distance), L2 (Euclidean distance), or  
1506 cosine distance. Specifically, for L1 and L2 distances,  $s_{ij} = \|x^{(i)} - x^{(j)}\|_p$  where  $p = 1$  for L1 and  
1507  $p = 2$  for L2, and for cosine distance,  $s_{ij} = 1 - \text{cosine\_similarity}(x^{(i)}, x^{(j)})$ .  
1508  
1509

1510 We denote by  $\hat{S} \in \mathbb{R}^{n \times n}$  the similarity matrix constructed from predicted responses, and by  $S \in$   
1511  $\mathbb{R}^{n \times n}$  the matrix for ground truth responses, with entries  $\hat{s}_{ij}$  and  $s_{ij}$ , respectively. The matrices are  
compared using the Frobenius norm.

#### 1512 K.4.1 BASIC SMD

1513 We define the Similarity Matrix Distance by comparing the similarity matrices using the Frobenius  
1514 norm:

$$1515 \text{SMD}(\hat{S}, S) = \|\hat{S} - S\|_F = \sqrt{\sum_{i=1}^n \sum_{j=1}^n (\hat{s}_{ij} - s_{ij})^2}, \quad (18)$$

1516 where  $\hat{s}_{ij}$  and  $s_{ij}$  are the similarity values between perturbations  $i$  and  $j$  in the predicted and ground  
1517 truth matrices, respectively. Lower values indicate better performance, as they represent smaller  
1518 differences between predicted and ground truth similarity structures. Expression data from all  
1519 compared models and baselines were normalized in the same way: total-count normalization to 10k  
1520 counts per cell followed by  $\log(1 + x)$ .

#### 1525 K.5 WEIGHTED MEAN SQUARED ERROR (WMSE)

1526 Following the calibrated metrics framework (Mejia et al., 2025), we evaluate perturbation predictions  
1527 using Weighted Mean Squared Error (WMSE), a weighted variant of the standard MSE metric  
1528 computed on pseudobulked data. For a single perturbation, WMSE is defined as:

$$1530 \text{WMSE} = \sum_{i=1}^G w_i (\mu_p^i - \hat{\mu}_p^i)^2, \quad (19)$$

1531 where  $G$  is the number of genes,  $w_i$  is the weight for gene  $i$ ,  $\mu_p^i$  is the observed mean expression for  
1532 gene  $i$  under perturbation  $p$ , and  $\hat{\mu}_p^i$  is the predicted mean expression for gene  $i$  under perturbation  $p$ .  
1533 Gene-specific weights  $w_i$  can be assigned to emphasize particular genes of interest (see Appendix G  
1534 for weight computation details).

1535 WMSE addresses limitations of unweighted error metrics by incorporating gene-level impor-  
1536 tance (Mejia et al., 2025). The weighting mechanism enables focus on genes that exhibit substantial  
1537 changes in response to perturbations, such as differentially expressed genes, rather than treating all  
1538 genes equally. This selective emphasis is crucial because most genes remain relatively unchanged  
1539 following perturbations, making it important to prioritize evaluation on genes with meaningful re-  
1540 sponses. Beyond evaluation, WMSE serves as a viable training objective that can replace standard  
1541 MSE, providing models with supervision that reflects biological relevance. Lower values indicate  
1542 better performance, corresponding to reduced weighted prediction errors.

#### 1546 K.6 MEAN SQUARED ERROR (MSE)

1547 MSE is a special case of WMSE K.5 with uniform weights  $w_i = 1/G$  for all genes:

$$1550 \text{MSE} = \frac{1}{G} \sum_{i=1}^G (\mu_p^i - \hat{\mu}_p^i)^2. \quad (20)$$

1551 Lower values indicate better performance.

#### 1555 K.7 MAXIMUM MEAN DISCREPANCY (MMD)

1556 Maximum Mean Discrepancy (MMD) measures the distributional similarity between predicted and  
1557 observed perturbed cell populations (Bunne et al., 2023). This metric is used to evaluate how well  
1558 the model captures the full distribution of cellular responses to perturbations, rather than just point  
1559 estimates or mean behavior. By comparing the distributions of predicted and observed cells, MMD  
1560 provides a comprehensive assessment of whether the model preserves the heterogeneity and variability  
1561 present in real perturbation responses.

1562 The mathematical formulation of MMD is provided in subsection J.1. We follow the approach  
1563 from (Bunne et al., 2023), using the energy kernel which corresponds to the Energy Distance metric.  
1564 Compared to positive-definite kernels such as the RBF, the energy kernel requires no bandwidth  
1565 hyperparameters.

## L BASELINES

### L.1 PERTURBATION PREDICTION

#### L.1.1 POST-PERTURBED MEAN

The post-perturbed mean baseline provides a dataset-wide reference by aggregating mean expression profiles across all training perturbations (Miller et al., 2025). For  $M$  training perturbations with cell sets  $S_1, S_2, \dots, S_M$  of potentially different sizes, we compute:

$$\mu_{\text{all}} = \frac{1}{M} \sum_{i=1}^M \left( \frac{1}{|S_i|} \sum_{j \in S_i} \mathbf{x}_{i,j} \right), \quad (21)$$

where  $\mathbf{x}_{i,j}$  denotes the expression vector of cell  $j$  from perturbation  $i$ . This formulation applies uniform weighting to each perturbation’s mean profile, regardless of the number of cells per perturbation. The resulting baseline represents a global average that smooths over perturbation-specific effects and cell-to-cell variability, serving as a simple but informative lower bound for model performance.

#### L.1.2 TECHNICAL DUPLICATE BASELINE

The technical duplicate baseline establishes an upper performance bound by leveraging the inherent variability within experimental replicates (Miller et al., 2025). For each perturbation  $p_i$  with associated cells  $S_i$ , we randomly partition the cells into two groups:  $S_{i,\text{GT}}$  serves as ground truth, while  $S_{i,\text{TD}}$  functions as a technical duplicate prediction. Mean expression profiles  $\mu_{i,\text{GT}}$  and  $\mu_{i,\text{TD}}$  are computed for each group and compared. This approach isolates the contribution of experimental noise to prediction error, effectively asking how well one experimental replicate predicts another under identical conditions. The baseline’s reliability depends on having sufficient cells per perturbation; with few cells, the random split introduces substantial variance. To maintain evaluation integrity, technical duplicate cells are treated as an independent held-out set and excluded from all differential expression analyses used for metric weighting.

#### L.1.3 INTERPOLATED DUPLICATE BASELINE

The interpolated duplicate baseline addresses the technical duplicate baseline’s limitations for perturbations with weak effects by adaptively blending predictions based on gene-level evidence of perturbation impact (Miller et al., 2025). For each perturbation, we construct a per-gene weighted combination:

$$\mu_{i,\text{ID}} = \alpha \odot \mu_{i,\text{TD}} + (1 - \alpha) \odot \mu_{\text{all}}, \quad (22)$$

where  $\alpha$  is a gene-specific weight vector and  $\odot$  represents element-wise multiplication. The weights are derived from differential expression analysis on the technical duplicate set:  $\alpha = 1 - p_{\text{DEGS}}$ , where  $p_{\text{DEGS}}$  contains  $p$ -values for each gene. Genes showing strong perturbation effects receive weights near 1, prioritizing the technical duplicate prediction, while genes with no significant changes receive weights near 0, defaulting to the mean baseline. This adaptive strategy provides a more realistic performance estimate by acknowledging that technical duplicates are most informative for genes with clear perturbation signatures.

## M ADDITIONAL IMPLEMENTATION DETAILS

This section provides additional implementation details on the task-specific model heads.

Task-specific heads process the compressed latent states  $L \in \mathbb{R}^{N_{\text{latent}} \times d}$  from the backbone to produce task-specific outputs. Each head architecture is tailored to the specific prediction task.

**Perturbation Prediction Head** For perturbation prediction, the head outputs gene expression vectors  $\hat{X} \in \mathbb{R}^{B \times G}$ . The latent states are projected back to gene space using a cross-attention layer that produces gene-level hidden states  $H_{\text{genes}} \in \mathbb{R}^{G \times d}$ . These gene-level representations are then mapped to expression predictions via a linear projection:  $\hat{X} = H_{\text{genes}} W_{\text{out}} + b_{\text{out}}$ , where  $W_{\text{out}} \in \mathbb{R}^{d \times G}$  and  $b_{\text{out}} \in \mathbb{R}^G$  are learnable parameters.

1620 **Unmasking Head** For unmasking tasks, the head architecture is similar to perturbation prediction,  
1621 outputting expression values for all genes including masked positions. The head processes latent  
1622 states to reconstruct the full expression profile:  $\hat{X} = \text{Head}_{\text{unmask}}(L) \in \mathbb{R}^{B \times G}$ .  
1623

1624 **Annotation Head (Classification)** For annotation tasks (cell type and tissue classification), the  
1625 head outputs class probabilities. The latent states are aggregated to  $\tilde{h} \in \mathbb{R}^d$ , then projected through a  
1626 linear layer to produce logits:  $z = \tilde{h}W_{\text{cls}} + b_{\text{cls}} \in \mathbb{R}^C$ , where  $C$  is the number of classes. The final  
1627 predictions are obtained via softmax:  $\hat{y} = \text{softmax}(z) \in [0, 1]^C$ .  
1628  
1629  
1630  
1631  
1632  
1633  
1634  
1635  
1636  
1637  
1638  
1639  
1640  
1641  
1642  
1643  
1644  
1645  
1646  
1647  
1648  
1649  
1650  
1651  
1652  
1653  
1654  
1655  
1656  
1657  
1658  
1659  
1660  
1661  
1662  
1663  
1664  
1665  
1666  
1667  
1668  
1669  
1670  
1671  
1672  
1673