Continual Causal Refinement: Learning from Sequential Perturbation Data

Anonymous Authors[1]

Abstract

Biological foundation models are often trained on large observational datasets collected across genomes, cells, tissues, or conditions, yet many downstream uses require accurate prediction under targeted perturbations. Perturbation data provide the relevant supervision, but each experiment is expensive and localized, producing small batches rather than broad resampling of the pretraining distribution. Refining a model from these batches therefore creates a continual learning problem: new perturbation evidence must be consolidated without erasing broad pretrained competence, and what is learned locally should transfer to future perturbation settings. We formalize this regime as continual causal refinement and instantiate it in regulatory genomics with a controlled benchmark in which sequence-to-function models are sequentially updated using in silico perturbation libraries. Benchmarking standard continual learning methods shows that naive fine-tuning causes substantial forgetting, whereas replay preserves observational performance and improves transfer to held-out perturbation libraries. Continual causal refinement provides a new avenue to iteratively refine biological foundation models with perturbation data.

1. Introduction

Foundation models are typically pretrained on large observational corpora and then refined through a sequence of post-training stages (Bommasani et al., 2021). The post-training stream is rarely more of the same data. It consists of curated batches that differ from the pretraining distribution in a structured way: human preference comparisons chosen for instructive feedback (Ouyang et al., 2022; Bai et al., 2022), robot trajectories sampled by an exploration policy (Kahn et al., 2015), or biological assays that intervene on the system being modeled (Kinney & McCandlish, 2019; Gasperini et al., 2019; Replogle et al., 2022). Observational pretraining recovers correlations supported by the data; the rules governing the underlying input space remain underdetermined, and targeted perturbation batches expose evidence that observational resampling cannot recover.

Standard task-incremental continual learning (CL) treats such streams as separate tasks and focuses on retaining each (Van de Ven & Tolias, 2019; De Lange et al., 2021). This does not match post-training, where batches sample the same underlying function under different perturbation strategies and the desired output is a single globally consistent model that generalizes to unseen interventional contexts.

We refer to this regime as *continual causal refinement*. The data are interventional in the statistical sense: each perturbation batch is generated by an intervention on the system being modeled, and the goal is to refine the model's approximation of the rules that govern that system. The continual aspect is intrinsic. Perturbations are typically localized and expensive, so adaptation must integrate evidence efficiently, which makes the order and selection of perturbation batches part of the problem itself.

Our contributions are: (i) a formal description of continual causal refinement and contrast with task-incremental CL; (ii) an instantiation in regulatory genomics with sequential in silico perturbation libraries; and (iii) a controlled empirical study with standard CL baselines (naive fine-tuning, replay, EWC) that exposes a retention and forward-transfer tension task-incremental evaluation does not capture.

2. Continual Causal Refinement

Standard task-incremental CL. A task-incremental CL stream is a sequence of tasks $t = 1, \ldots, T$, each with its own distribution $P_t(X, Y)$ and possibly its own label space Y_t. The objective combines task losses, $\min_\theta \sum_{t=1}^{T} \mathbb{E}_{(x,y)\sim P_t}[\ell(f_\theta(x), y)]$, under memory and access constraints, focusing on minimizing forgetting of earlier tasks (Goodfellow et al., 2013; Van de Ven & Tolias, 2019; De Lange et al., 2021).

Continual causal refinement setup. We distinguish the biological response from the experimental measurement.

[1]Anonymous Institution, Anonymous City, Anonymous Region, Anonymous Country. Correspondence to: <anon@domain.com>.

Preliminary work. Under review by the ICML 2026 Workshop "Continual Adaptation at Scale: Towards Sustainable AI". Do not distribute.

Let $g : \mathcal{X} \to \mathcal{Z}$ denote the biological input-response map, where $\mathcal{X}$ is the input space and $\mathcal{Z}$ is the latent biological state space. For $x \in \mathcal{X}$, $z = g(x) \in \mathcal{Z}$ is the induced latent state. Let $\mathcal{E}$ be the set of experiment types. Each $e \in \mathcal{E}$ observes the latent response through an assay-specific map $h_e : \mathcal{Z} \to \mathcal{Y}_e$, giving observable target $f_e^\star = h_e \circ g$. Labels satisfy $y = h_e(g(x)) + \varepsilon_e$, with assay-dependent noise ε_e. Pretraining data and perturbation batches are

$$\begin{aligned} \mathcal{D}_{\text{obs}} &= \{(x_i^0, e_i^0, y_i^0)\}_{i=1}^{n_0}, \quad x_i^0 \sim P_{\text{obs}}, \\ y_i^0 &= h_{e_i^0}(g(x_i^0)) + \varepsilon_i^0, \\ \mathcal{D}_t &= \{(x_i^t, e_t, y_i^t)\}_{i=1}^{n_t}, \quad x_i^t \sim Q_t, \\ y_i^t &= h_{e_t}(g(x_i^t)) + \varepsilon_i^t. \end{aligned} \tag{1}$$

where $e_t \in \mathcal{E}$ is the experiment type at step t, and Q_t is the input distribution induced by the perturbation strategy.

The continual learner updates $\theta_t = A(\theta_{t-1}, \mathcal{D}_t, M_t)$, where M_t denotes auxiliary memory such as replay data, regularization statistics, or modular adapters. We write the predictor as $f_\theta(x, e)$, with $f_\theta(\cdot, e) : \mathcal{X} \to \mathcal{Y}_e$, allowing the prediction map to depend on the experiment type.

The objective is global observable-function recovery:

$$\min_{\theta_T} \ \mathbb{E}_{(x,e)\sim P_{\text{eval}}} \left[\ell_e\big(f_{\theta_T}(x, e), f_e^\star(x)\big) \right], \tag{2}$$

where P_{eval} may put mass on the observational marginal, past interventional contexts $Q_1, \ldots, Q_T$, and future contexts $Q_{T+1}, \ldots$ not observed during training.

Terminology. We use "causal" in the mechanistic sense common to physics and biology: the rules that generate observed activity and remain stable under intervention. In this paper, the relevant interventions are perturbations of the biological system, not formal SCM interventions (Pearl, 2000; Schölkopf et al., 2021). SCM machinery is not invoked because no identifiability arguments are made.

Distinguishing properties. Continual causal refinement differs from task-incremental CL in three ways. First, all batches are noisy samples of the same underlying mechanism g, so the CL stream is over input distributions and experiment-specific readouts h_{e_t}; there is no per-task label space. Second, the pretraining marginal P_{obs} is consistent with many mechanisms in the regions it undersamples, and the interventional marginals Q_t expose regions that distinguish among them, carrying information observational resampling cannot recover. Third, success is measured by the quality of the global function approximation, including extrapolation to unseen interventional contexts; task retention is a special case restricted to past distributions.

Relation to adjacent regimes. Each component of continual causal refinement has been studied separately, but their combination in scientific foundation-model post-training has no established name. Covariate shift and domain adaptation share $P(Y|X)$ across distributions and address one-shot transfer (Shimodaira, 2000; Ben-David et al., 2010); sequential foundation-model adaptation falls outside their scope. Continual domain adaptation handles sequential shifts, but those shifts are often natural or adversarial drift. Active learning queries a pool to reduce labels (Ren et al., 2021); it is rarely paired with foundation-model updating and is not structurally interventional. Sequential Bayesian experimental design and active scientific machine learning come closest (Shahriari et al., 2015; Stanton et al., 2022), though they typically use small bespoke models. Invariance-based methods such as IRM share the goal of recovering predictors that hold under perturbation (Arjovsky et al., 2019), but assume access to multiple environments during training. Continual causal adaptation describes how these elements combine when a foundation model pretrained on observational data is refined through sequential targeted perturbations on a shared system.

Evaluation criteria. Three evaluation criteria follow. *Retention* measures whether performance on the original observational distribution is preserved after adaptation. *Backward transfer* measures whether perturbation batches improve observational performance beyond the pretraining baseline, thereby implying retention; *forward transfer* measures generalization to unseen interventional contexts (Lopez-Paz & Ranzato, 2017).

Scientific foundation models with the same regime. The continual causal refinement pattern recurs across scientific foundation models. In single-cell biology, transcriptomic foundation models pretrained on observational scRNA-seq atlase are refined with Perturb-seq data combining CRISPR perturbations with single-cell readouts (Cui et al., 2024). In protein modeling, structural foundation models pretrained on natural PDB structures are refined with deep mutational scanning and directed-evolution data that intervene on individual residues (Lafita et al., 2024). In materials discovery, models pretrained on existing crystal databases are refined through autonomous-lab synthesis and characterization rounds, each round consuming wall-clock time and reagents (Merchant et al., 2023). We focus on regulatory genomics because the observational and interventional split is unusually clean, the foundation models are public, and locality is well-defined.

3. A Controlled Testbed in Genomics

Sequence-to-function (seq2func) models map DNA sequence to regulatory readouts such as chromatin accessibility, transcription factor binding, or gene expression, and are typically pretrained on genome-wide functional genomics

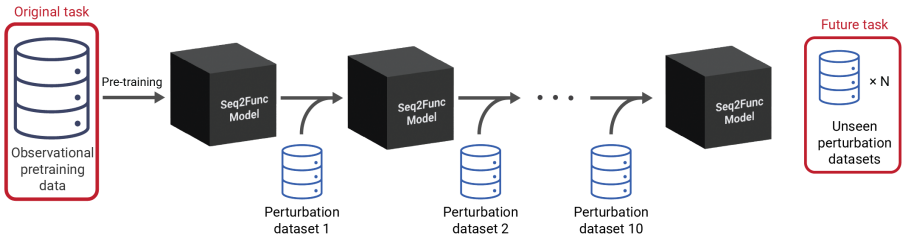

Figure 1. **Continual causal refinement pipeline.** A pretrained sequence-to-function model is sequentially updated on perturbation libraries arriving one at a time (datasets 1 to T). At each step, N held-out perturbation libraries serve as future interventional contexts, and performance on these measures forward transfer.

assays aligned to a reference genome (Pampari et al., 2025; Avsec et al., 2026). Biophysical and regulatory rules govern sequence-to-activity relationships across all of sequence space, but these rules are not known explicitly. The reference genome samples a confounded slice of this space and is consistent with many different rule sets, so genome-wide performance does not imply generalization to engineered variants or alternative regulatory configurations (Sasse et al., 2023; Huang et al., 2023; Karollus et al., 2023).

Perturbation assays provide the interventional counterpart. Massively parallel reporter assays (MPRAs) (Kinney & McCandlish, 2019) measure the activity of designed or mutated regulatory sequences. CRISPR interference (CRISPRi) screens measure how suppressing regulatory elements affects nearby gene expression. Both directly probe cis-regulatory mechanism, and they must do so locally: perturbing every regulatory element at once would disrupt cellular state, so individual assays target specific loci. Each MPRA or CRISPRi assay corresponds to one Q_t in the formulation above, a localized interventional distribution that probes part of the rule set; the continual problem is integrating these local probes efficiently into a model that must generalize across all of sequence space.

Why in silico. Wet-lab regulatory perturbation rounds can take months to years, making multi-step continual learning streams impractical for method development. We therefore use a stronger pretrained seq2func model as an oracle to label in silico perturbation libraries, enabling rapid prototyping of adaptation schedules, CL strategies, and evaluation protocols. Each experiment contains sequences derived from one locus through a fixed perturbation strategy (local sequence variation for MPRAs, element silencing for CRISPRi), each oracle-labeled, and each step of the stream provides one batch of perturbation data (Figure 1). This oracle landscape is only a controlled approximation: it inherits model biases and may differ from biological reality. Thus, conclusions about which CL method is best in vivo are tentative; the regime structure and evaluation protocol should transfer more reliably.

4. Empirical Study

Stream and evaluation protocol. Each step of the stream provides one perturbation batch $\mathcal{D}_t$. The student parameters θ_t are produced by updating θ_{t-1} on $\mathcal{D}_t$. We evaluate at every step on two held-out sets. *Retention* and *backward transfer* are both measured on the genome-wide observational test set used to pretrain the student: performance at the pretraining baseline indicates retention while performance that exceeds it indicates backward transfer. *Forward transfer* is measured on a fixed set of held-out perturbation libraries from loci that never enter the CL stream. All metrics are reported as Pearson correlation between predictions and oracle labels.

Baselines. Because continual causal refinement is a new benchmark setting, we use simple reference baselines to establish the problem. We compare two well-established CL methods, *replay* and *elastic weight consolidation* (EWC), and include naive fine-tuning as a baseline. Naive fine-tuning updates the student only on the perturbation batch $\mathcal{D}_t$. Replay mixes samples from $\mathcal{D}_{\text{obs}}$ into each adaptation step using $\mathcal{L} = \mathcal{L}_{\text{int}} + \alpha\mathcal{L}_{\text{replay}}$, with fixed mixing ratio α (Rolnick et al., 2019). EWC penalizes movement in parameters identified as important for pretrained performance using $\mathcal{L} = \mathcal{L}_{\text{int}} + \frac{\lambda}{2}\sum_i F_i(\theta_i - \theta_i^*)^2$, where the diagonal Fisher F is estimated once on $\mathcal{D}_{\text{obs}}$ at the start of the stream and the anchor θ^* is set to θ_0 (Kirkpatrick et al., 2017).

4.1. Specialist Student with in silico MPRAs

Setup. The student is ProCapNet (Cochran et al., 2024), which predicts transcription initiation profiles and counts from ∼2kb DNA sequence. The oracle is MPRALegNet (Agarwal et al., 2025). Each MPRA experiment is constructed by selecting a high-activity sequence and generating local sequence variants, which is labeled by the oracle. The stream contains ten libraries from distinct loci, and the forward-transfer set consists of MPRA on ten additional loci that never enter the stream.

Results. Naive fine-tuning underwent severe forgetting, as expected. Profile correlation drops from ≈ 0.5 to near zero

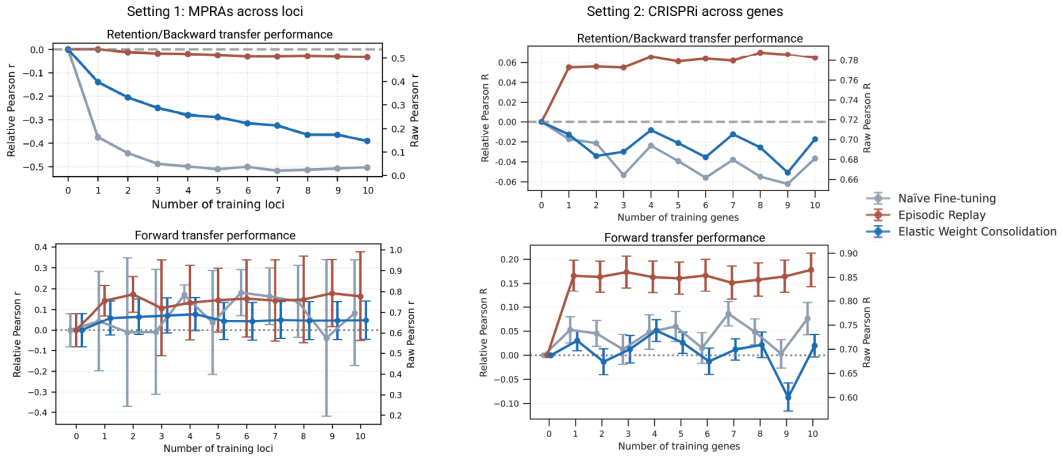

Figure 2. **Genome-wide retention and forward transfer across 10 sequential adaptation steps.** Columns show Setting 1 (ProCapNet + MPRAs, left) and Setting 2 (Enformer + CRISPRi, right). *Top row:* Pearson r on the held-out genome-wide test set relative to the pretrained baseline after each adaptation step; values below zero indicate forgetting, above zero indicate backward transfer. *Bottom row:* Forward transfer to held-out perturbation libraries not seen during adaptation: 10 MPRA libraries (left) and 50 CRISPRi libraries (right). Naive fine-tuning (gray), replay (red), EWC (blue).

(Figure 2, top left). Replay preserves the original genome-wide performance and improves forward transfer to held-out MPRA experiments by roughly 0.2 in Pearson correlation (bottom left). On the other hand, EWC yields poor performance, slowly decaying toward the naive trajectory across the refinement rounds.

4.2. Generalist Student with in silico CRISPRi

Setup. The student is Enformer (Avsec et al., 2021), a multitask seq2func model that predicts thousands of functional genomics data from long context DNA sequences. The oracle is Decima (Lal et al., 2024). Each CRISPRi experiment targets a single gene and consists of sequences in which one or more candidate regulatory elements are silenced in silico (Karollus et al., 2023; Toneyan & Koo, 2024); labels are given by Decima's predicted gene expression. The stream contains ten CRISPRi experiments on different target genes. Forward transfer was evaluated using CRISPRi experiments across fifty held-out genes. Before sequential refinement, we calibrate the student to the oracle on unperturbed genomic sequences, so that adaptation is measured relative to a shared baseline prediction scale.

Results. Naive fine-tuning reduced performance on the observational test set from ≈ 0.72 to 0.67 Pearson correlation (Figure 2, top right), indicating forgetting. Replay improved performance to ≈ 0.78, consistent with backward transfer, while EWC remained near baseline at ≈ 0.70. For forward transfer, replay was the only method that improved performance on held-out perturbation libraries, increasing Pearson correlation by roughly 0.17 (Figure 2, bottom right). Together, these results show that replay can both preserve observational performance and improve generalization to unseen perturbation contexts in this refinement regime.

5. Conclusion

Many biological foundation models are pretrained on large observational datasets, but their intended use often requires prediction under perturbation. Perturbation data provide the relevant evidence, yet each experiment samples only a localized part of the input space. Continual causal refinement frames this as a CL problem: sequential perturbation batches should improve a shared model without erasing broad observational competence, while transferring what is learned to future perturbation settings. Our results show that this setting is informative even with simple baselines. Naive fine-tuning forgets the observational distribution, whereas replay preserves observational performance and improves forward transfer to held-out perturbation libraries. In the CRISPRi setting, replay also improves observational performance, suggesting that perturbation data can sharpen the broader predictor when consolidated appropriately.

The key takeaway is that continual causal refinement exposes the retention-transfer tradeoff in perturbation-based refinement. This raises several open questions: what should be replayed, how replay and parameter-efficient fine-tuning (PEFT) (Hu et al., 2022; Houlsby et al., 2019) should coordinate, and how the next perturbation batch should be chosen (See Appendix A.1 for a discussion of why PEFT is complementary but not sufficient). Because observational training leaves many input-output rules underdetermined, perturbation selection should prioritize interventions that distinguish among plausible model explanations and expose dependencies that transfer beyond the assayed context.

Impact Statement

This work studies continual adaptation methodology in a controlled in silico setting and makes no biological or clinical claims.

References

Agarwal, V., Inoue, F., Schubach, M., Penzar, D., Martin, B. K., Dash, P. M., Keukeleire, P., Zhang, Z., Sohota, A., Zhao, J., et al. Massively parallel characterization of transcriptional regulatory elements. *Nature*, 639(8054): 411–420, 2025.

Arjovsky, M., Bottou, L., Gulrajani, I., and Lopez-Paz, D. Invariant risk minimization. *arXiv preprint arXiv:1907.02893*, 2019.

Avsec, Ž., Agarwal, V., Visentin, D., Ledsam, J. R., Grabska-Barwinska, A., Taylor, K. R., Assael, Y., Jumper, J., Kohli, P., and Kelley, D. R. Effective gene expression prediction from sequence by integrating long-range interactions. *Nature Methods*, 18(10):1196–1203, 2021.

Avsec, Ž., Latysheva, N., Cheng, J., Novati, G., Taylor, K. R., Ward, T., Bycroft, C., Nicolaisen, L., Arvaniti, E., Pan, J., et al. Advancing regulatory variant effect prediction with alphagenome. *Nature*, 649(8099):1206–1218, 2026.

Bai, Y., Jones, A., Ndousse, K., Askell, A., Chen, A., DasSarma, N., Drain, D., Fort, S., Ganguli, D., Henighan, T., et al. Training a helpful and harmless assistant with reinforcement learning from human feedback. *arXiv preprint arXiv:2204.05862*, 2022.

Ben-David, S., Blitzer, J., Crammer, K., Kulesza, A., Pereira, F., and Vaughan, J. W. A theory of learning from different domains. *Machine learning*, 79(1):151–175, 2010.

Bommasani, R., Hudson, D. A., Adeli, E., Altman, R., Arora, S., von Arx, S., Bernstein, M. S., Bohg, J., Bosselut, A., Brunskill, E., et al. On the opportunities and risks of foundation models. *arXiv preprint arXiv:2108.07258*, 2021.

Cochran, K., Yin, M., Mantripragada, A., Schreiber, J., Marinov, G. K., Shah, S. R., Yu, H., T LIS, J., and Kundaje, A. Dissecting the cis-regulatory syntax of transcription initiation with deep learning. *BioRxiv*, 2024.

Cui, H., Wang, C., Maan, H., Pang, K., Luo, F., Duan, N., and Wang, B. scgpt: toward building a foundation model for single-cell multi-omics using generative ai. *Nature methods*, 21(8):1470–1480, 2024.

De Lange, M., Aljundi, R., Masana, M., Parisot, S., Jia, X., Leonardis, A., Slabaugh, G., and Tuytelaars, T. A continual learning survey: Defying forgetting in classification tasks. *IEEE transactions on pattern analysis and machine intelligence*, 44(7):3366–3385, 2021.

Gasperini, M., Hill, A. J., McFaline-Figueroa, J. L., Martin, B., Kim, S., Zhang, M. D., Jackson, D., Leith, A., Schreiber, J., Noble, W. S., et al. A genome-wide framework for mapping gene regulation via cellular genetic screens. *Cell*, 176(1):377–390, 2019.

Goodfellow, I. J., Mirza, M., Xiao, D., Courville, A., and Bengio, Y. An empirical investigation of catastrophic forgetting in gradient-based neural networks. *arXiv preprint arXiv:1312.6211*, 2013.

Houlsby, N., Giurgiu, A., Jastrzebski, S., Morrone, B., De Laroussilhe, Q., Gesmundo, A., Attariyan, M., and Gelly, S. Parameter-efficient transfer learning for nlp. In *International conference on machine learning*, pp. 2790–2799. PMLR, 2019.

Hu, E. J., Shen, Y., Wallis, P., Allen-Zhu, Z., Li, Y., Wang, S., Wang, L., Chen, W., et al. Lora: Low-rank adaptation of large language models. *Iclr*, 1(2):3, 2022.

Huang, C., Shuai, R. W., Baokar, P., Chung, R., Rastogi, R., Kathail, P., and Ioannidis, N. M. Personal transcriptome variation is poorly explained by current genomic deep learning models. *Nature Genetics*, 55(12):2056–2059, 2023.

Kahn, G., Sujan, P., Patil, S., Bopardikar, S., Ryde, J., Goldberg, K., and Abbeel, P. Active exploration using trajectory optimization for robotic grasping in the presence of occlusions. In *2015 IEEE International Conference on Robotics and Automation (ICRA)*, pp. 4783–4790. IEEE, 2015.

Karollus, A., Mauermeier, T., and Gagneur, J. Current sequence-based models capture gene expression determinants in promoters but mostly ignore distal enhancers. *Genome biology*, 24(1):56, 2023.

Kinney, J. B. and McCandlish, D. M. Massively parallel assays and quantitative sequence–function relationships. *Annual Review of Genomics and Human Genetics*, 20(1): 99–127, 2019.

Kirkpatrick, J., Pascanu, R., Rabinowitz, N., Veness, J., Desjardins, G., Rusu, A. A., Milan, K., Quan, J., Ramalho, T., Grabska-Barwinska, A., et al. Overcoming catastrophic forgetting in neural networks. *Proceedings of the National Academy of Sciences*, 114(13):3521–3526, 2017.

Lafita, A., Gonzalez, F., Hossam, M., Smyth, P., Deasy, J., Allyn-Feuer, A., Seaton, D., and Young, S. Fine-tuning protein language models with deep mutational scanning improves variant effect prediction. *arXiv preprint arXiv:2405.06729*, 2024.

Lal, A., Karollus, A., Gunsalus, L., Garfield, D., Nair, S., Tseng, A. M., Gordon, M. G., Blischak, J., van de Geijn, B., Bhangale, T., et al. Decoding sequence determinants of gene expression in diverse cellular and disease states. *bioRxiv*, pp. 2024–10, 2024.

Lopez-Paz, D. and Ranzato, M. Gradient episodic memory for continual learning. *Advances in neural information processing systems*, 30, 2017.

Merchant, A., Batzner, S., Schoenholz, S. S., Aykol, M., Cheon, G., and Cubuk, E. D. Scaling deep learning for materials discovery. *Nature*, 624(7990):80–85, 2023.

Ouyang, L., Wu, J., Jiang, X., Almeida, D., Wainwright, C., Mishkin, P., Zhang, C., Agarwal, S., Slama, K., Ray, A., et al. Training language models to follow instructions with human feedback. *Advances in neural information processing systems*, 35:27730–27744, 2022.

Pampari, A., Shcherbina, A., Kvon, E. Z., Kosicki, M., Nair, S., Kundu, S., Kathiria, A. S., Risca, V. I., Kuningas, K., Alasoo, K., et al. ChromBPNet: bias factorized, base-resolution deep learning models of chromatin accessibility reveal cis-regulatory sequence syntax, transcription factor footprints and regulatory variants. *bioRxiv*, 2025.

Pearl, J. *Causality: Models, Reasoning and Inference*. Cambridge University Press, New York, 2000.

Ren, P., Xiao, Y., Chang, X., Huang, P.-Y., Li, Z., Gupta, B. B., Chen, X., and Wang, X. A survey of deep active learning. *ACM computing surveys (CSUR)*, 54(9):1–40, 2021.

Replogle, J. M., Saunders, R. A., Pogson, A. N., Hussmann, J. A., Lenail, A., Guna, A., Mascibroda, L., Wagner, E. J., Adelman, K., Lithwick-Yanai, G., et al. Mapping information-rich genotype-phenotype landscapes with genome-scale perturb-seq. *Cell*, 185(14):2559–2575, 2022.

Rolnick, D., Ahuja, A., Schwarz, J., Lillicrap, T., and Wayne, G. Experience replay for continual learning. *Advances in Neural Information Processing Systems*, 32, 2019.

Sasse, A., Ng, B., Spiro, A. E., Tasaki, S., Bennett, D. A., Gaiteri, C., De Jager, P. L., Chikina, M., and Mostafavi, S. Benchmarking of deep neural networks for predicting personal gene expression from dna sequence highlights shortcomings. *Nature genetics*, 55(12):2060–2064, 2023.

Schölkopf, B., Locatello, F., Bauer, S., Ke, N. R., Kalchbrenner, N., Goyal, A., and Bengio, Y. Toward causal representation learning. *Proceedings of the IEEE*, 109(5): 612–634, 2021.

Shahriari, B., Swersky, K., Wang, Z., Adams, R. P., and De Freitas, N. Taking the human out of the loop: A review of bayesian optimization. *Proceedings of the IEEE*, 104 (1):148–175, 2015.

Shimodaira, H. Improving predictive inference under covariate shift by weighting the log-likelihood function. *Journal of statistical planning and inference*, 90(2):227–244, 2000.

Stanton, S., Maddox, W., Gruver, N., Maffettone, P., Delaney, E., Greenside, P., and Wilson, A. G. Accelerating bayesian optimization for biological sequence design with denoising autoencoders. In *International conference on machine learning*, pp. 20459–20478. PMLR, 2022.

Toneyan, S. and Koo, P. K. Interpreting cis-regulatory interactions from large-scale deep neural networks. *Nature genetics*, 56(11):2517–2527, 2024.

Van de Ven, G. M. and Tolias, A. S. Three scenarios for continual learning. *arXiv preprint arXiv:1904.07734*, 2019.

A. Appendix

A.1. Relation to PEFT

Parameter-efficient fine-tuning (PEFT) methods such as LoRA and adapters specify *where* updates are applied but not *how* evidence should be consolidated across interventional batches. The two concerns are largely orthogonal: PEFT reduces the parameter footprint of each adaptation step, while continual causal refinement governs what the updates should preserve, consolidate, and transfer.

This distinction matters in practice. Per-task adapters avoid interference across interventional batches but fragment the model into task-specific modules, which is at odds with the goal of a single globally consistent approximation of g. A single shared adapter avoids fragmentation but reintroduces interference, reproducing the same forgetting and forward-transfer tension as naive fine-tuning within a lower-dimensional parameter subspace. PEFT is therefore complementary to the methods studied here: replay, EWC, or other consolidation strategies can be applied over PEFT parameters just as over full model parameters, but PEFT alone does not resolve the continual causal refinement problem.