

Multimodal Latent Causal VAE for Joint Inference of Gene Regulatory and Protein Interaction Networks



Paper Link



Berkeley Engineering

Shervin Goudarzi^{1,2}, Nuo Liu¹, Hani Goodarzi^{1,3}

Arc Institute¹, UC Berkeley², UCSF³

Introduction

Learning directed causal relationships between genes and proteins from purely observational single-cell data remains a fundamental challenge. We introduce **CrossModal-CausalVAE (CMC-VAE)**, a latent-based causal VAE that learns directed graphs within and across modalities via linear decoders and a DAGMA-constrained structural causal model. Trained on 809K CITE-seq cells, the model recovers central dogma directionality: RNA→Protein edges are significantly stronger than reverse for matched pairs alongside prior-consistent (i.e. CollecTRI) transcription factor→target asymmetry. **Zero-shot validation** against held-out Perturb-CITE-seq data achieves RNA→Protein perturbation AUROC of 0.739 and 69% gene-level direction accuracy.

Data

Dataset	Cells	Donors	Condition
Hao, 2021	161,764	8	Healthy PBMC
Stephenson, 2021	647,366	130	COVID-19 + Healthy
Total	809,130	138	Mixed

Model Data	Prior Knowledge
80% / 10% / 10%	GRN: CollecTRI (OmniPath) — 1,604 gene interactions
Random split by cell	PPI: OmniPath directed interactions — 144 protein pairs
2000 HVGs + 144 Proteins	Batch correction: donor-level (138 donors)

$$\underbrace{\mathcal{L}_{\text{rec}}}_{\text{reconstruction}} + \beta \underbrace{D_{\text{KL}}}_{\text{KL div.}} + \lambda_{\text{dag}} \underbrace{h(A)}_{\text{acyclicity}} + \lambda_{\text{prior}} \underbrace{\mathcal{L}_{\text{margin}}}_{\text{bio. prior}} + \lambda_{\text{sparse}} \underbrace{\|A\|_1}_{\text{graph sp.}} + \lambda_{\text{dec}} \underbrace{\|W\|_1}_{\text{dec. sp.}} + \lambda_{\text{ortho}} \underbrace{\|W^T W - I\|_F^2}_{\text{orthogonality}}$$

Equation 1: Loss Function

Results

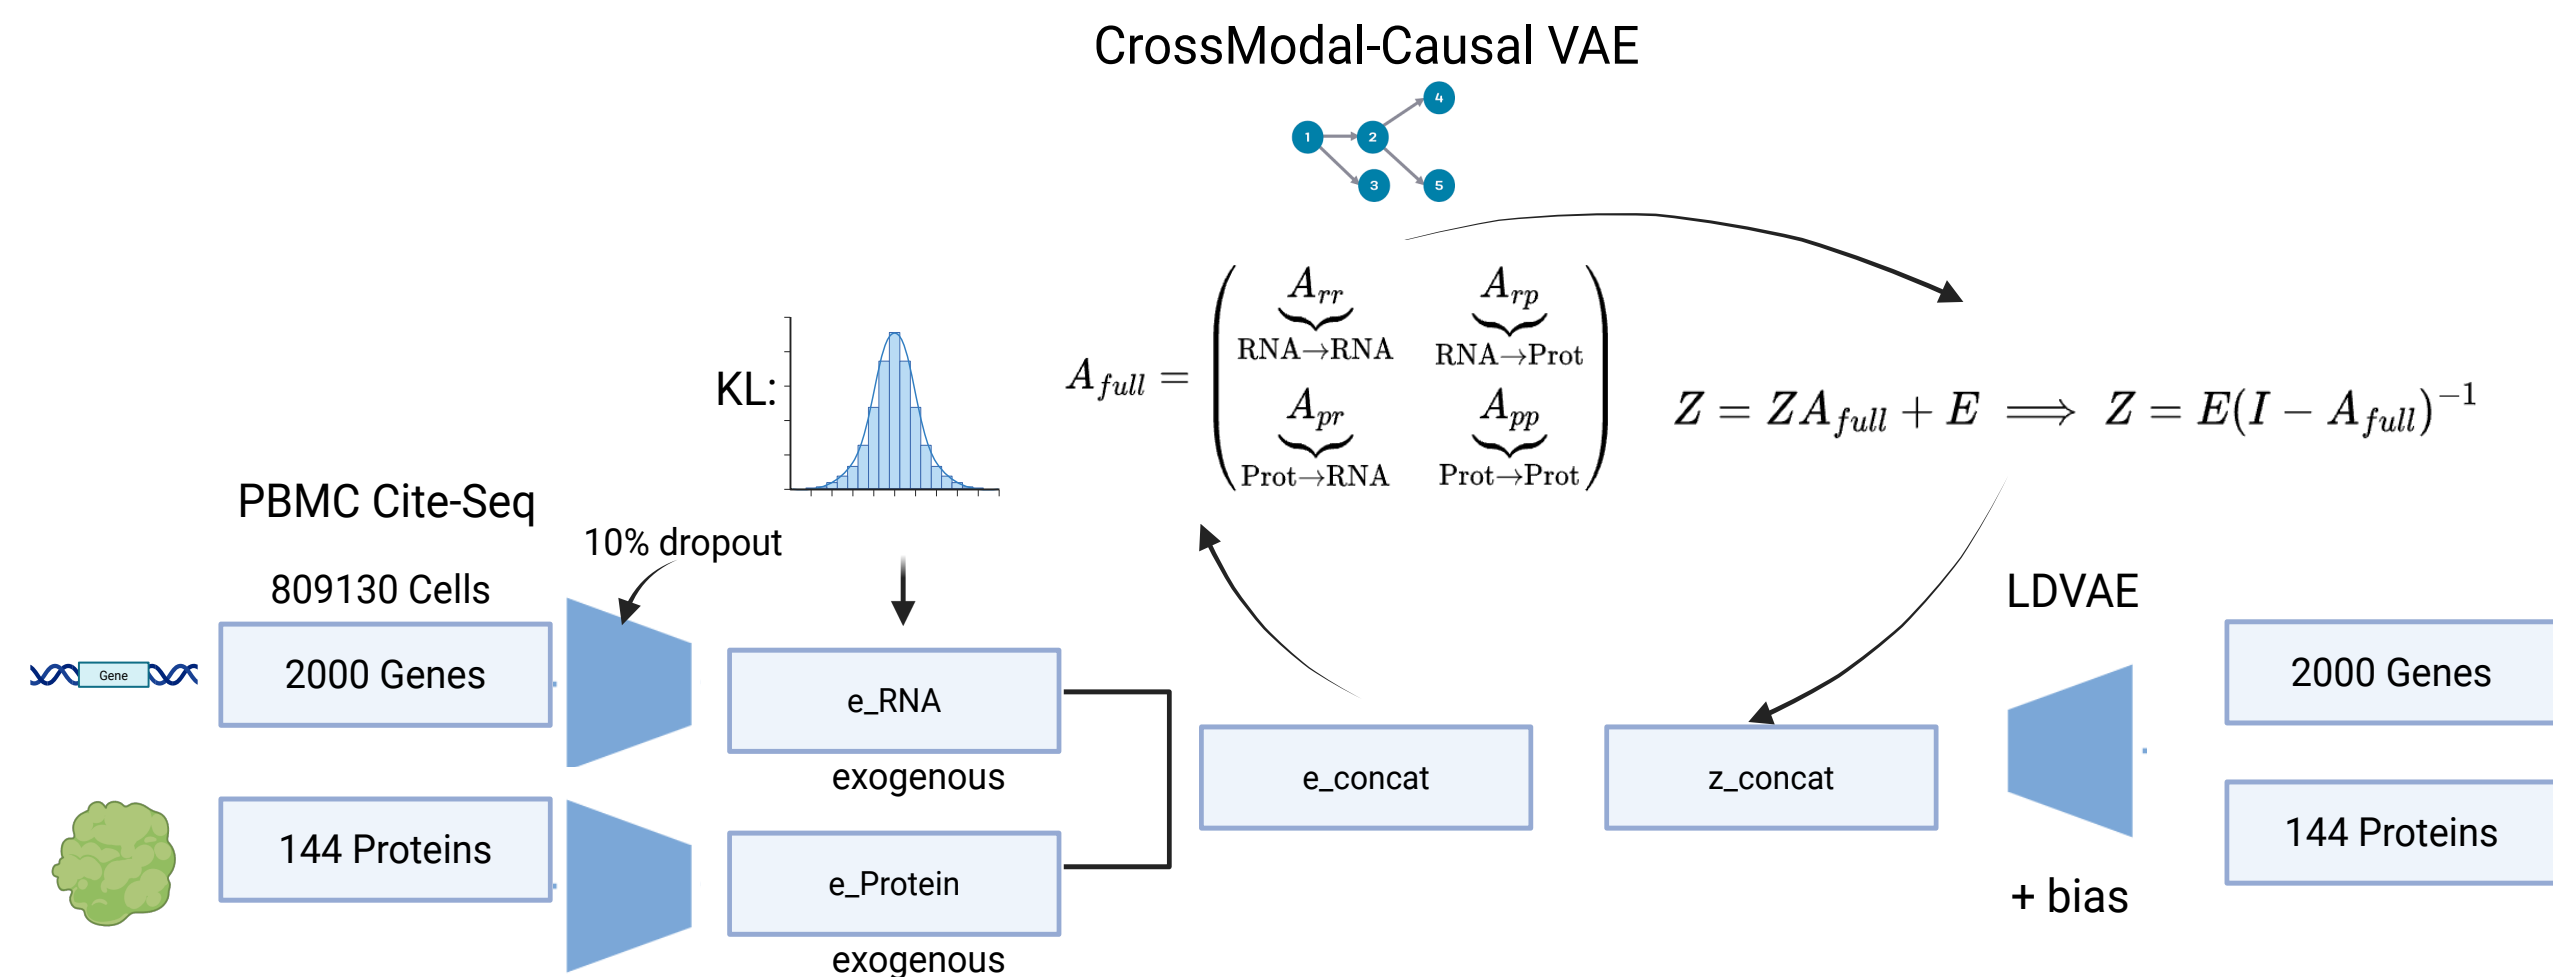


Figure 1: CMC-VAE Model Architecture

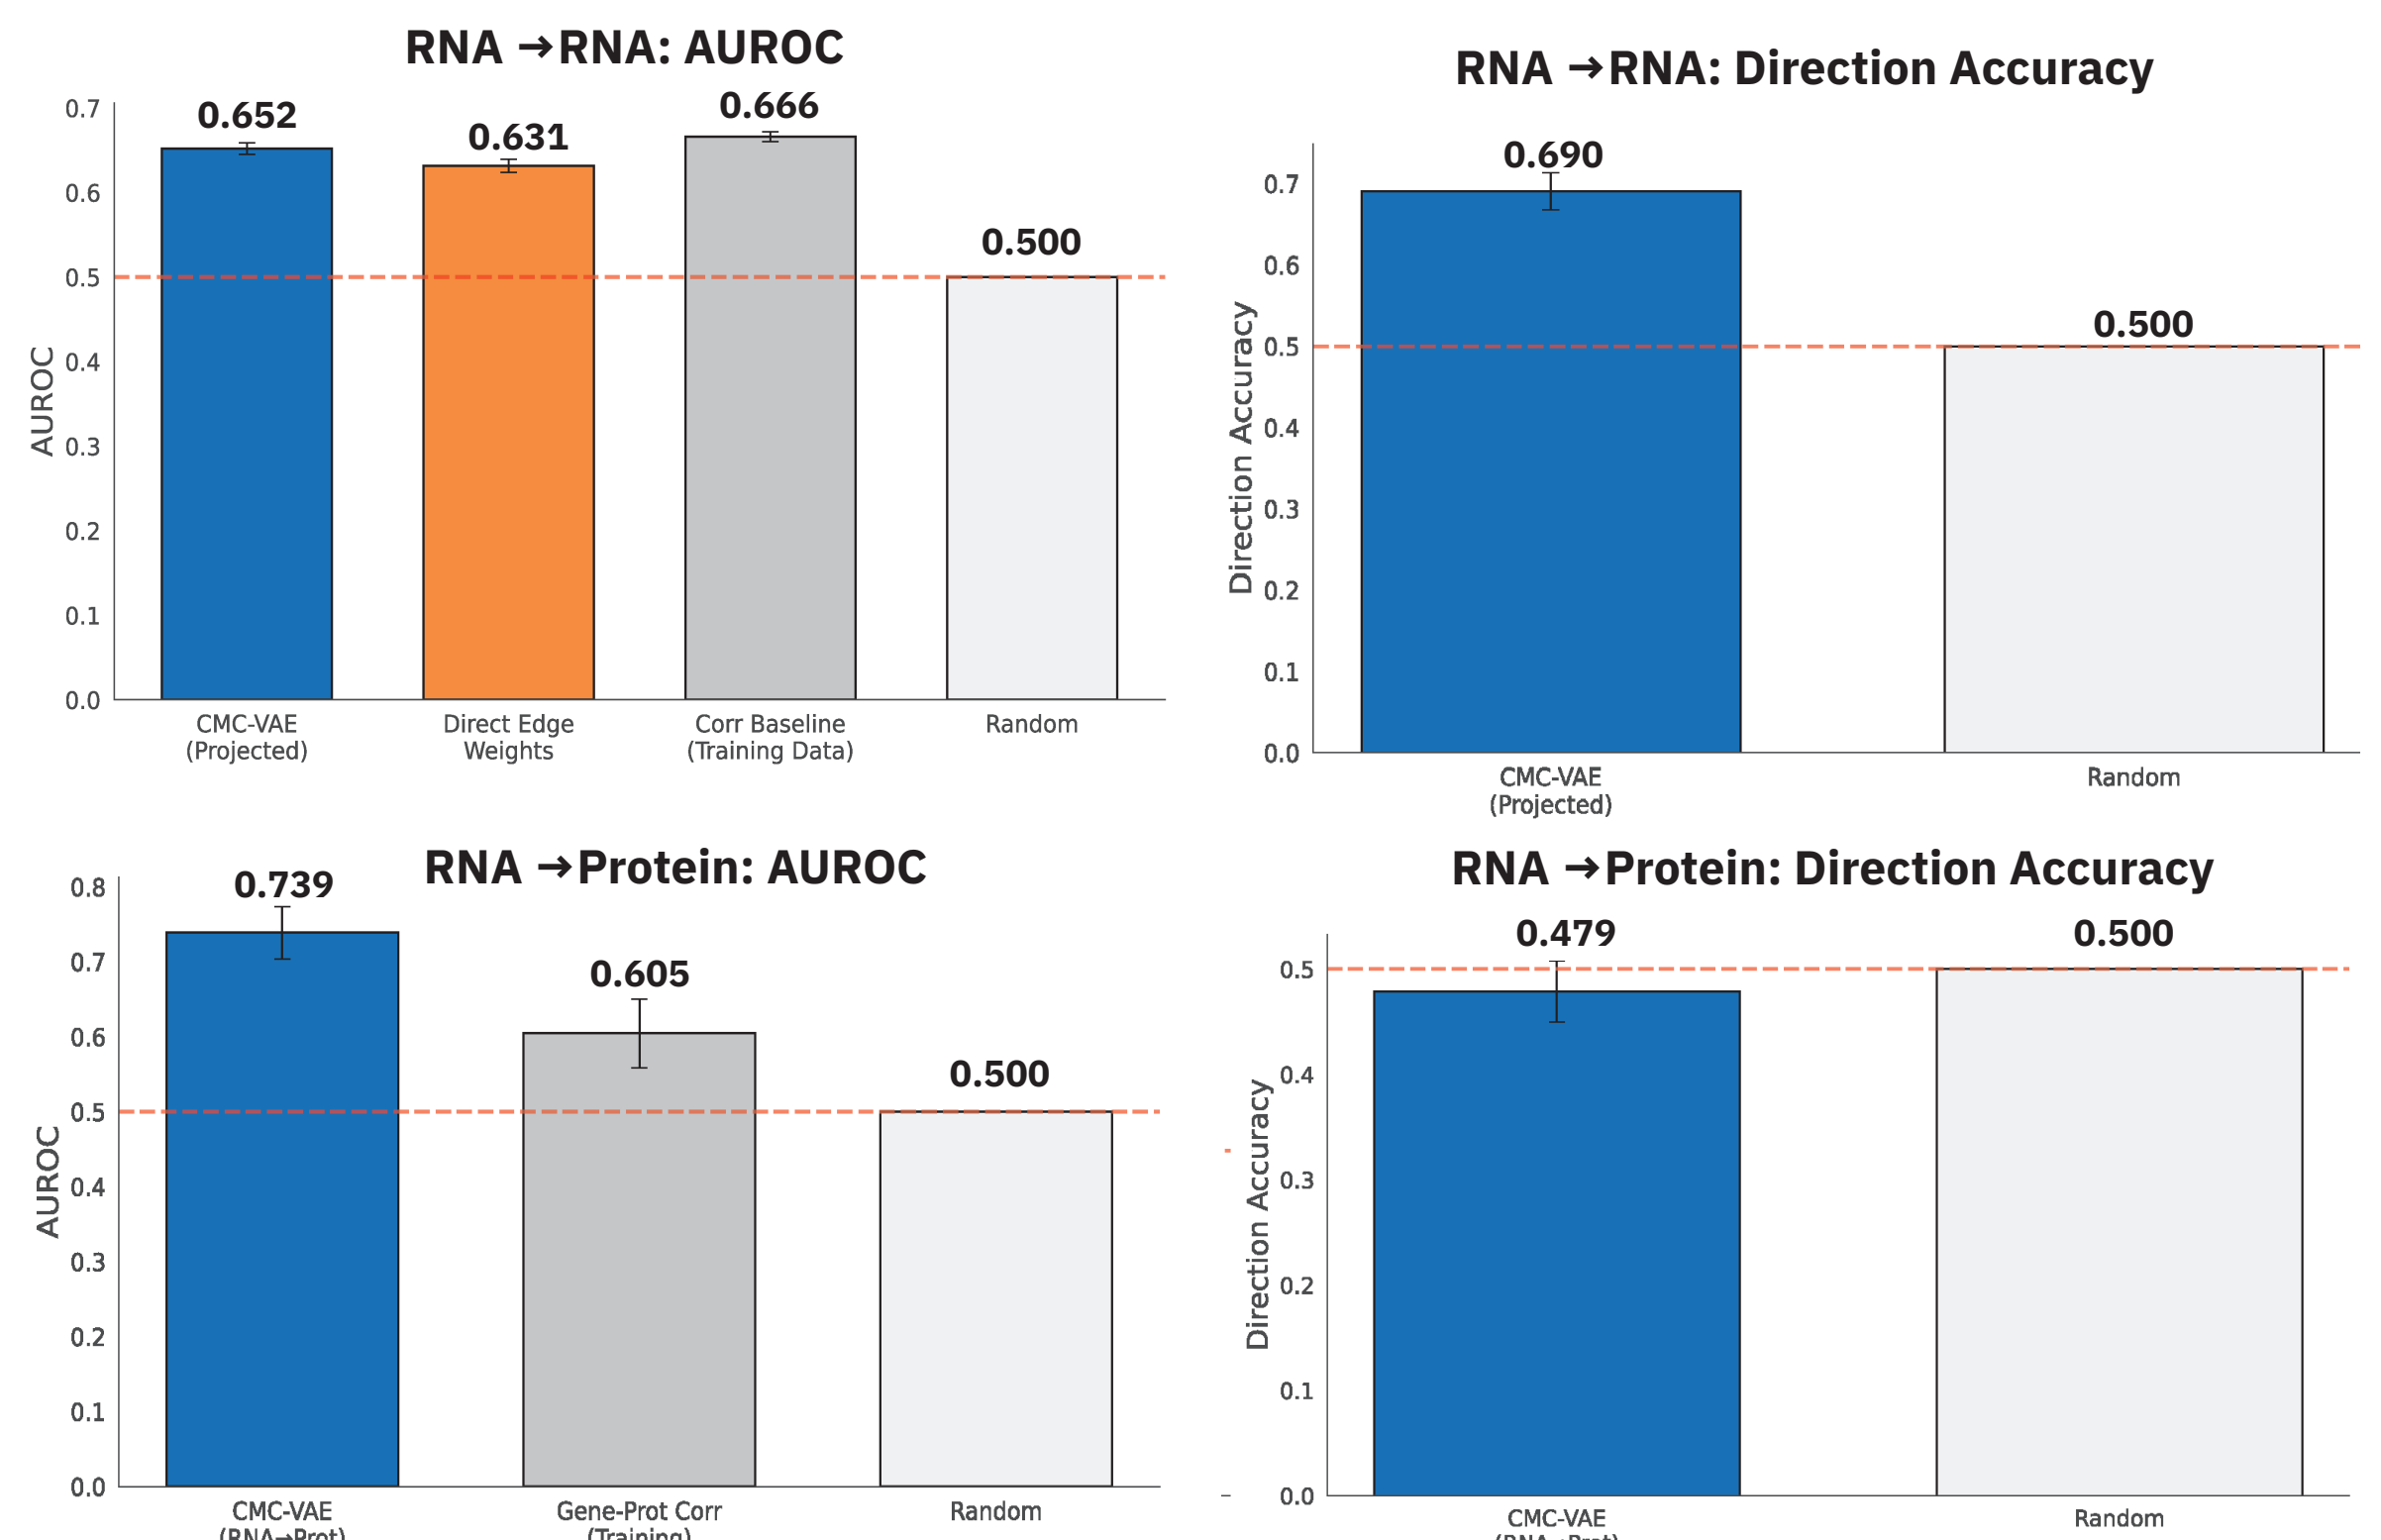


Figure 2: Zero-shot Perturb Cite-Seq Validation 37 Perturbations




Figure 3: Regulatory Asymmetry

Figure 4: Central Dogma Asymmetry

Conclusion

CMC-VAE learns meaningful **causal structure** from purely **observational multi-modal** single-cell data. By embedding a structural causal model in the VAE latent space with linear decoders, the model recovers biologically grounded, directed relationships, validated against known priors and **held-out perturbation data**.

Hao, Y., Hao, S., Andersen-Nissen, E., et al. "Integrated analysis of multimodal single-cell data." Cell, 184(13), 3573–3587. (2021).

Stephenson, E., Reynolds, G., Botting, R.A., et al. "Single-cell multi-omics analysis of the immune response in COVID-19." Nature Medicine, 27, 904–916. (2021).

Frangieh, C.J., Melms, J.C., Thakore, P.I., et al. "Multimodal pooled Perturb-CITE-seq screens in patient models define mechanisms of cancer immune evasion." Nature Genetics, 53, 332–341. (2021).

Türei, D., Korcsmáros, T., Sáez-Rodríguez, J. "OmniPath: guidelines and gateway for literature-curated signaling pathway resources." Nature Methods, 13, 966–967. (2016).

Yang, M., Liu, F., Chen, Z., et al. "CausalVAE: Disentangled representation learning via neural structural causal models." CVPR, 9593–9602. (2021).