SUPPLEMENTARY MATERIALS FOR:
SCHYPERLINK: REVEALING CELL-TYPE-SPECIFIC GENE REGULATION WITH HYPERGRAPH NEURAL NETWORKS

Emre Kulkul, Tolga Çukur, Aykut Koç

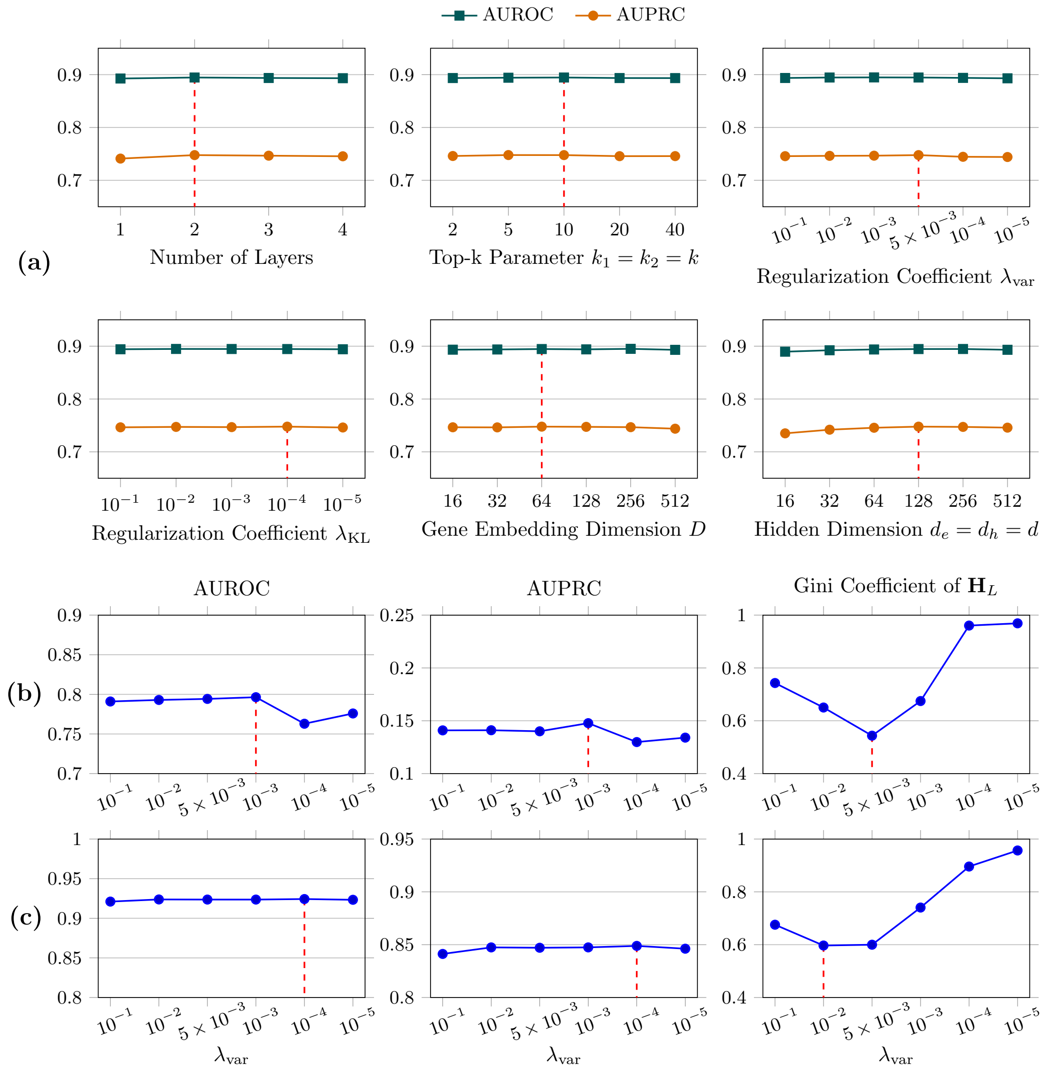

Supp. Fig. 1: (a) Validation AUROC and AUPRC scores obtained as an average across all cell types with a hyperparameter sweep. The best-performing hyperparameters (either in terms of AUROC or AUPRC, depending on the difference) are used as the experimental setup for scHyperLink. They are denoted here with red dashed lines. (b) AUROC, AUPRC, and Gini coefficient of the learned hypergraph matrix $\mathbf{H}_L$ versus λ_{var} on the validation set of the mDC-500 dataset, (c) and the mESC-1000 dataset. The best-performing point with respect to the corresponding metric is denoted with red dashed lines.

Supp. Table. 1: Summary of cell-type-specific scRNA-seq data after pre-processing with the most varying 500 (and $1,000$) genes.

Cell Type	Cells	All Genes	All TFs
hESC	759	910 (1,410)	410 (410)
hHEP	426	948 (1,448)	448 (449)
mDC	384	821 (1,321)	323 (323)
mESC	422	1,120 (1,620)	627 (628)
mHSC-E	1,072	704 (1,204)	205 (209)
mHSC-GM	890	632 (1,132)	135 (136)
mHSC-L	848	560 (692)	61 (61)

Supp. Table. 2: Running time and peak GPU VRAM usage of different baselines on the mHSC-L-500 dataset.

mHSC-L-500	DeepSEM	HyperG-VAE	GENELink	GNNLink	GATCL	GMFGRN	GRACE	scMGATGRN	GCLink	LineGRN	scHyperLink
Running Time (s)	47.65	48.26	31.21	2.55	25.72	85.22	12.21	15.61	22.33	840.92	33.92
Used VRAM (MB)	927	3,236	1,241	784	1,579	1,141	1,379	1,367	1,445	22,553	1,257

Supp. Table. 3: Running time and peak GPU VRAM usage of different baselines on the hHEP-1000 dataset.

hHEP-1000	DeepSEM	HyperG-VAE	GENELink	GNNLink	GATCL	GMFGRN	GRACE	scMGATGRN	GCLink	LineGRN	scHyperLink
Running Time (s)	82.56	52.30	148.76	6.91	110.72	518.53	15.02	53.88	110.00	5,728.57	166.68
Used VRAM (MB)	1,235	3,351	1,557	1,140	1,815	839	1,947	2,073	2,545	20,065	1,367

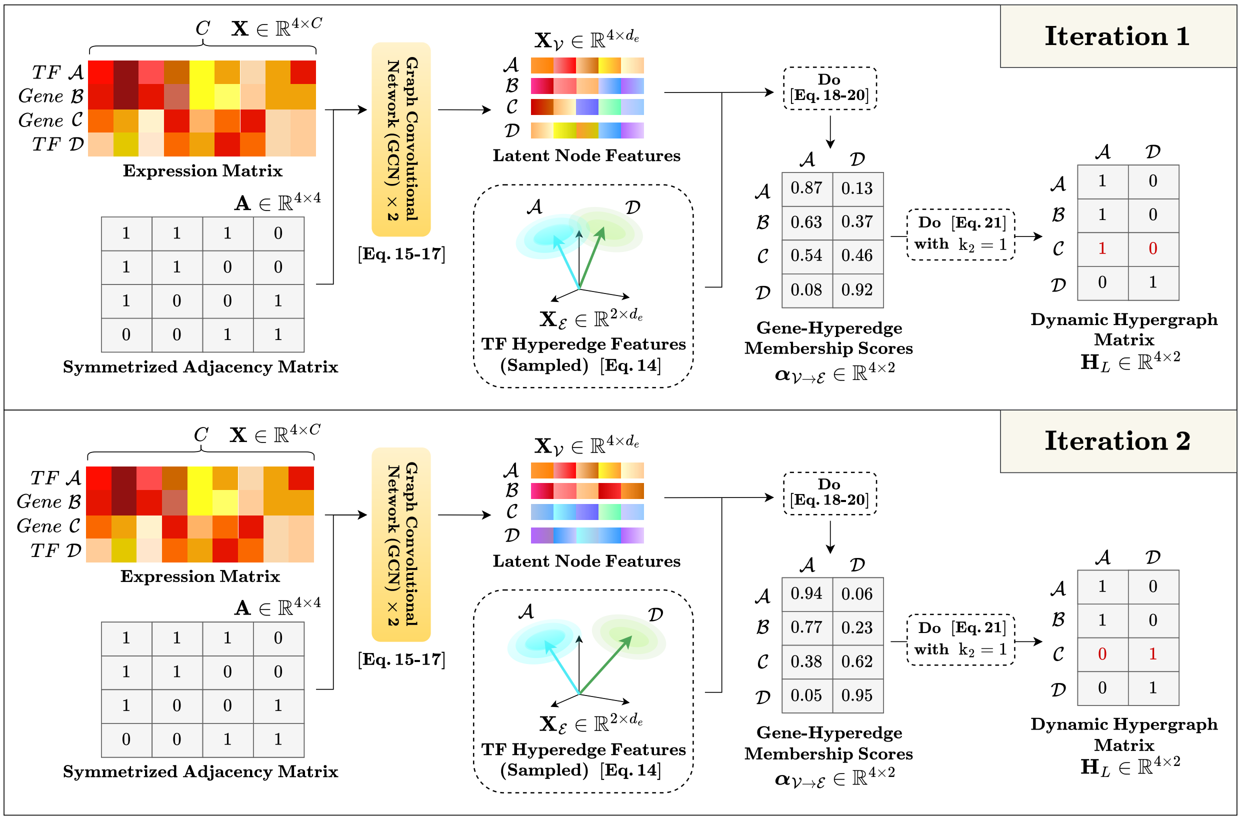

Supp. Fig. 2: A toy example to demonstrate the principles of the DHI module. The expression matrix $\mathbf{X}$ contains two TFs A and D and two target genes B and C, where (A, B) and (C, D) exhibit correlated expression. The prior adjacency $\mathbf{A}$ encodes known TF–target links but includes a spurious edge (A, C), which may arise from database noise or indirect co-expression artifacts captured in bulk-derived interaction resources. In the first iteration, this prior causes C to be incorrectly assigned to A's hyperedge. As training proceeds, dynamic updates to the latent node features $\mathbf{X}_v$ and TF prototypes $\mathbf{X}_e$ suppress the unreliable prior and reassign C to D's hyperedge, aligning memberships with true co-expression structure.

Supp. Table. 4: OSN co-targeting overlap vs. size-matched null distribution case study.

Statistic	**Value**
Observed mean OSN overlap	0.01635
Null mean ± std. dev.	0.00878 ± 0.0017
Exceedances (≥ observed)	5/10000
Empirical p-value	5.99×10^{-4}
Fold enrichment	1.863×