EviGRN

Multimodal Graph Reasoning for Gene Regulatory Network Inference

Akshata Hegde[1,2], Kyler Zook[1], Yanli Wang[1], and Jianlin Cheng[1,2]

[1] Department of Electrical Engineering and Computer Science
[2] Roy Blunt NextGen Precision Health, University of Missouri

Poster – ID: 149734, Dec 10 · Poster Session 03 · Hall C1

Mizzou

NeurIPS 2026 · Dec 9–13
Atlanta, Georgia

NEURAL INFORMATION PROCESSING SYSTEMS

Gene regulatory networks and their evidence

A gene regulatory network links transcription factors to their target genes.

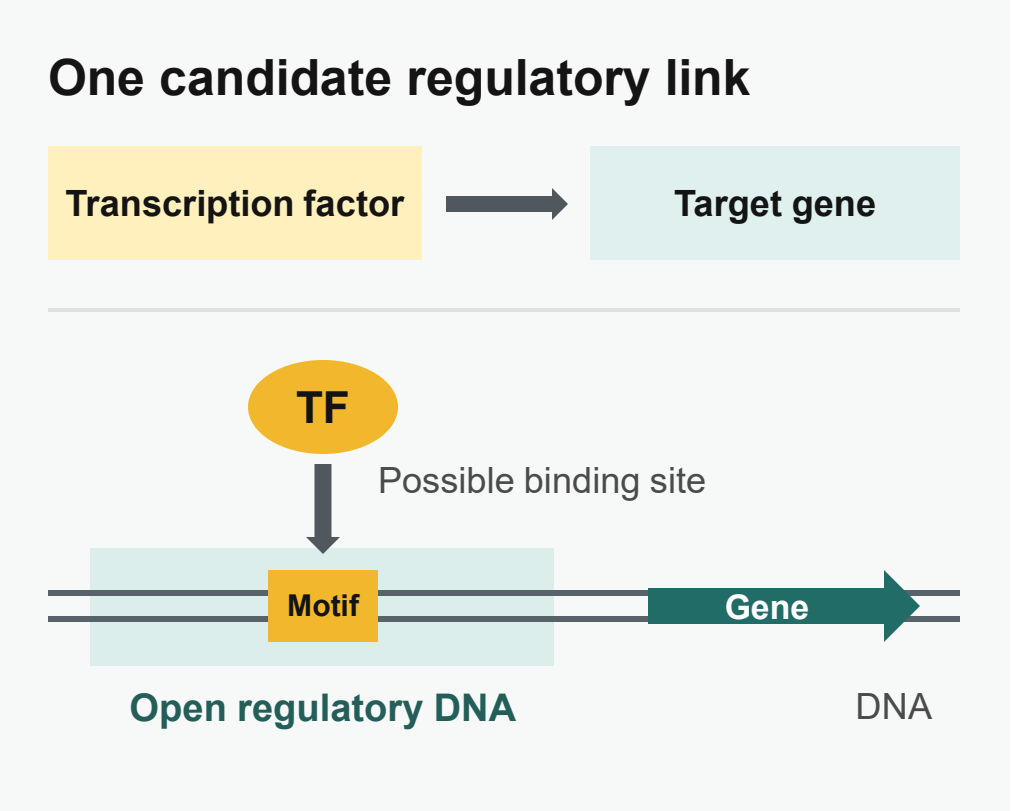

RNA expression
Which genes are active, and at what levels?

Chromatin accessibility
Which DNA regions are open to binding?

TF-binding motifs
Which DNA sequences could a regulator recognize?

Evidence is often incomplete.
Its relevance depends on the biological context.

EviGRN pipeline overview

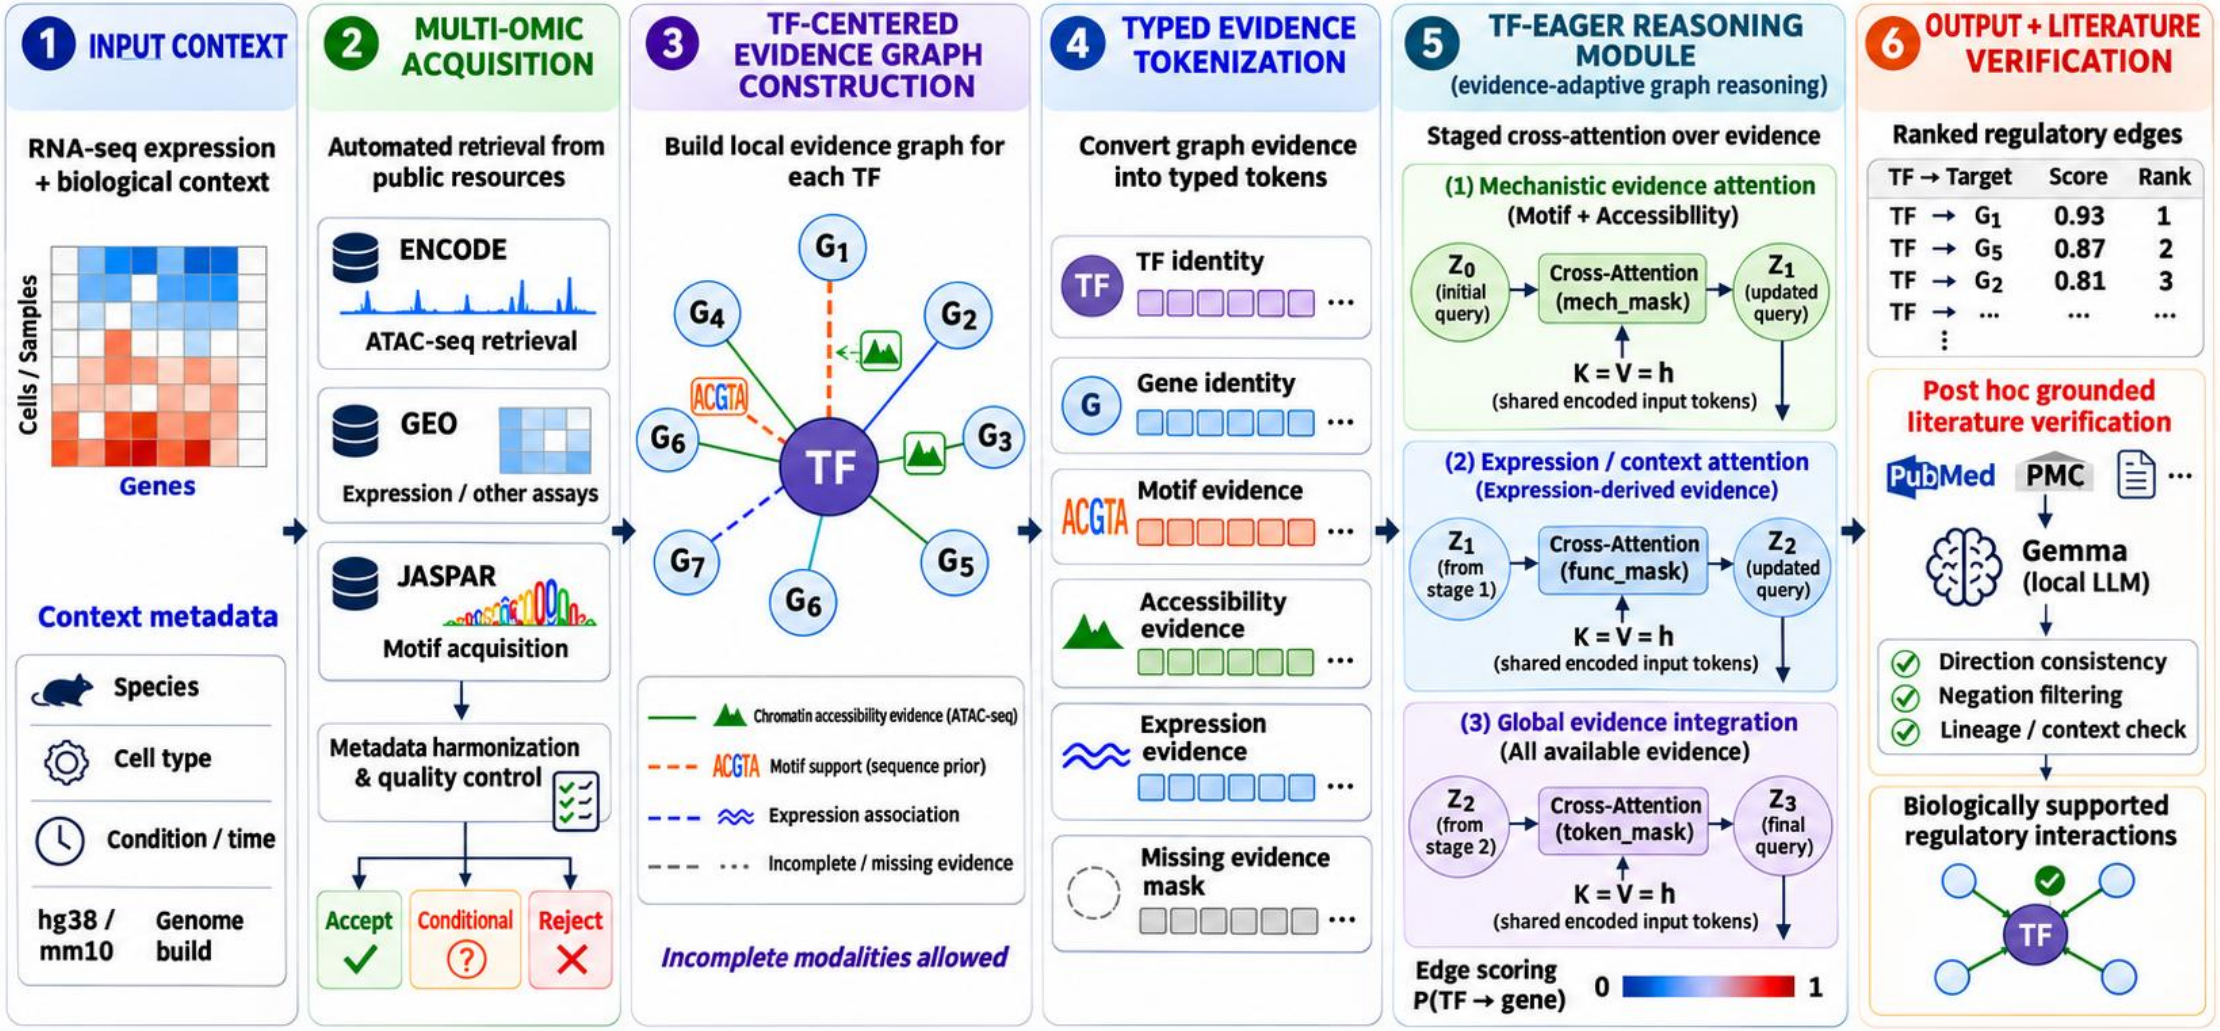

EviGRN: Multi-modal Evidence adaptive Graph Reasoning for GRN Inference

- Automated acquisition and QC of multi-omic evidence
- TF-centered evidence graph with incomplete modalities
- Typed evidence tokens with missingness modeling
- Staged evidence integration via TF-EAGER
- Robust inference under missing data and across contexts

TF	Transcription Factor
G	Target Gene
ATAC-seq	Assay for Transposase-Accessible Chromatin
LLM	Large Language Model

Held-out benchmark results

4.9× the mean AUPRC of the strongest baseline

0.0311 vs 0.0063 for GRNFormer across 24 full-matrix settings and ten baselines

Sampled evaluation across positive-to-negative ratios

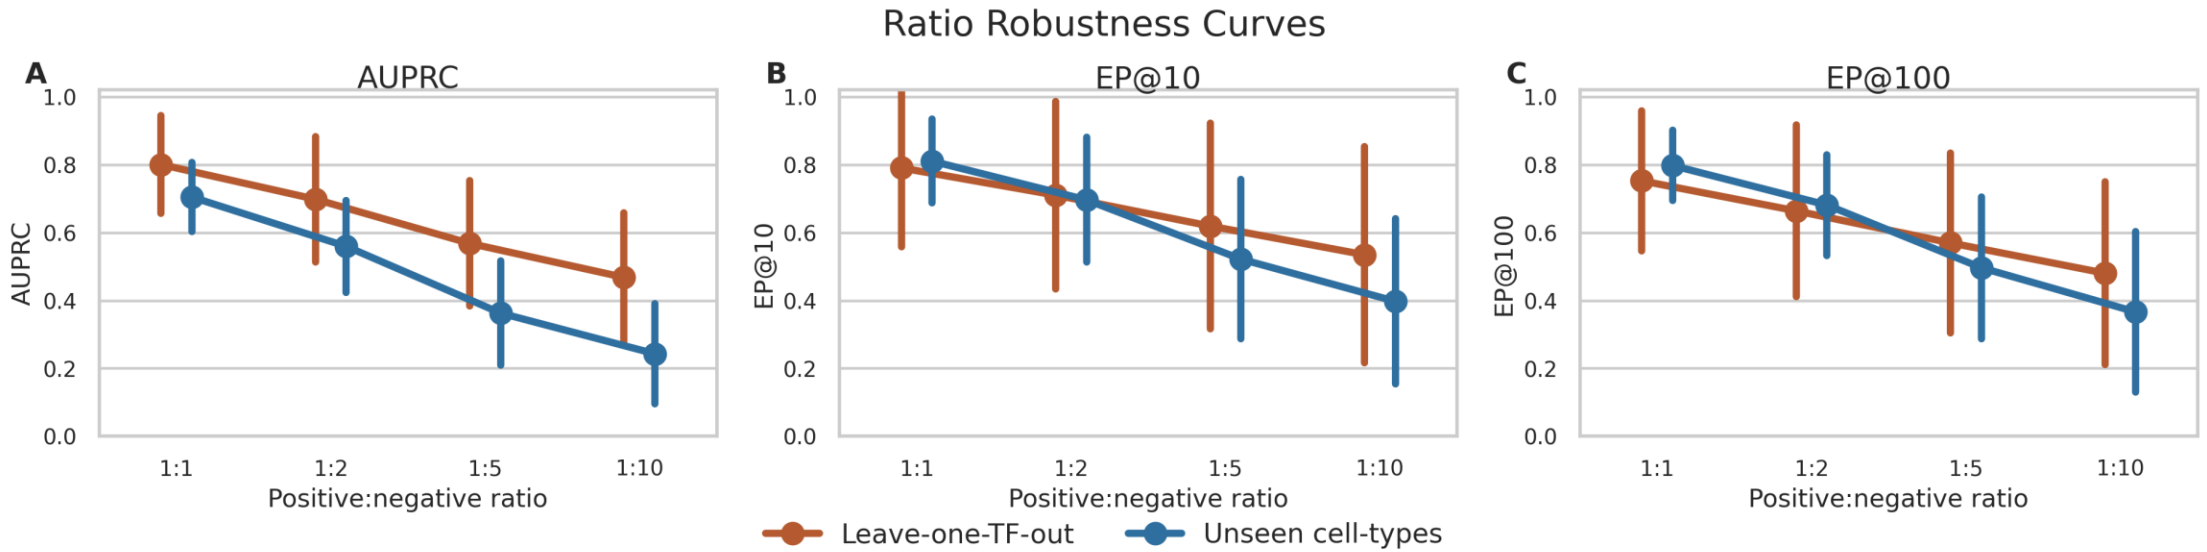

At 1:1 sampling, AUPRC was 0.801 for leave-one-TF-out and 0.705 for unseen cell types.

RNA-only baselines and added evidence

Mean AUPRC with RNA-only inputs

Reference network	EviGRN	Strongest RNA baseline
Cell-type-specific ChIP-seq	**0.084**	**0.016 (CNNC)**
Non-specific ChIP-seq	0.005	0.002 (GRNFormer)
STRING	0.004	0.002 (GRNFormer)

Removing motifs and accessibility before rebuilding the graph

38 settings with complete evidence

Mean AUPRC fell by 0.060. Mean early precision at 100 fell by 0.068.

Comparison with multimodal methods

Four human contexts: EviGRN on each baseline's own candidate interactions

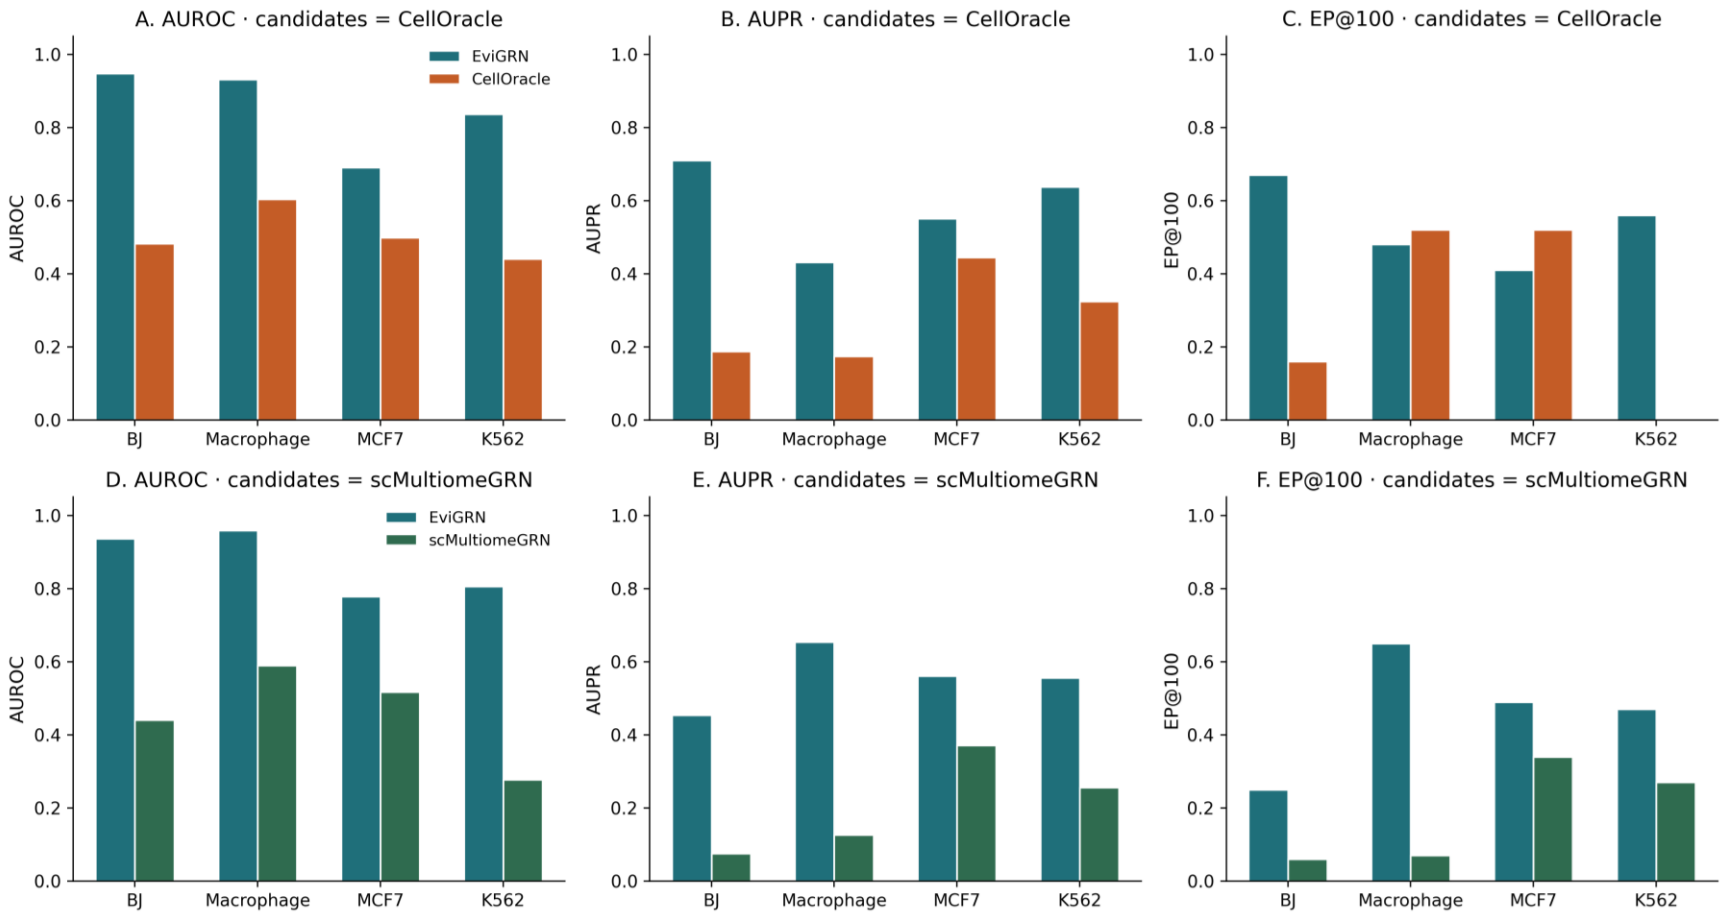

Mean AUPRC

CellOracle pool

0.583 vs 0.283

scMultiomeGRN pool

0.556 vs 0.207

EviGRN vs comparator

Biological case study and literature validation

Mouse myeloid progenitors

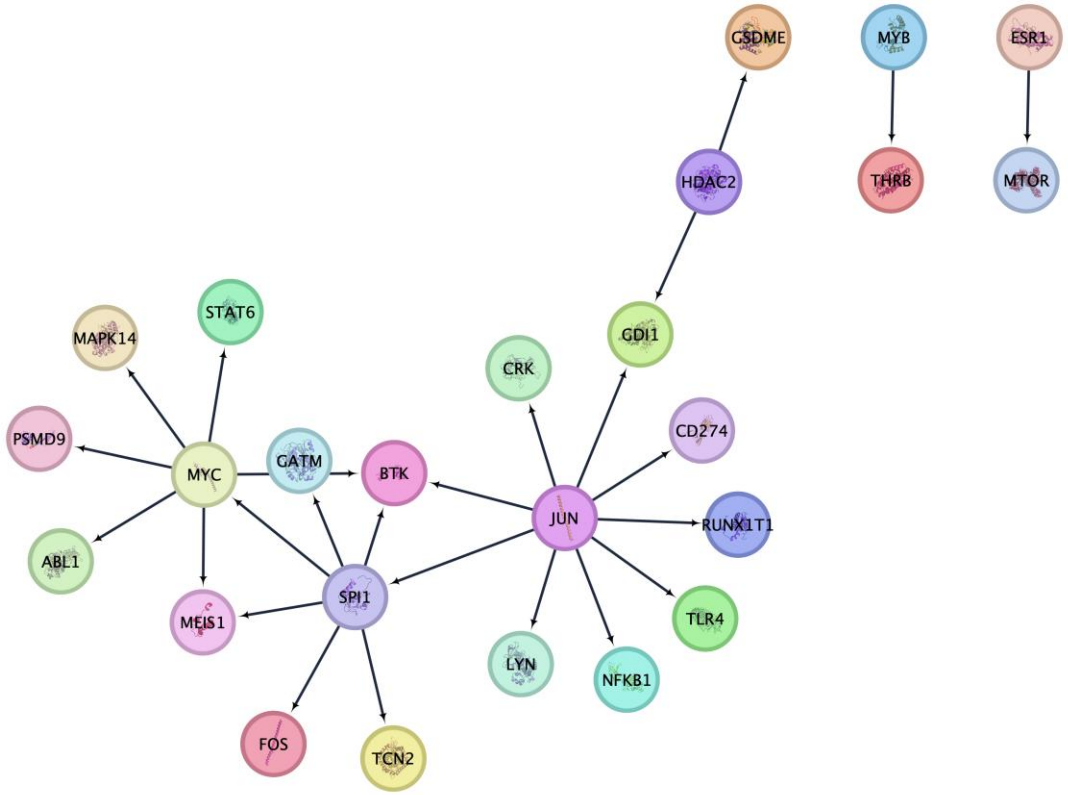

25 literature-supported examples from a broader prediction pool

Recovered known regulatory circuitry

Including circuits around SPI1, JUN, and MYB

Support beyond the benchmark

17 of 100

top-ranked edges absent from the benchmark had published support.

Literature verification controls

72%

of known interactions had literature support

50 of 50

fabricated pairs were rejected

Generalization across species

The mammalian-trained model transferred without retraining

E. coli

DREAM5 benchmark

Early precision of 0.60 for the top 10 predictions

At 1:1 positive-to-negative sampling

Drosophila

Eye-antennal disc, 3,531 cells

Predicted relationships linked to Notch and Hedgehog signaling

Qualitative analysis of developmental pathways

EviGRN interface and code

Run inference

Provide expression data, transcription factors and the biological context.

Explore predictions

Inspect candidate neighborhoods, evidence coverage and individual edge scores.

Code, pretrained weights and interface

github.com/BioinfoMachineLearning/EviGRN

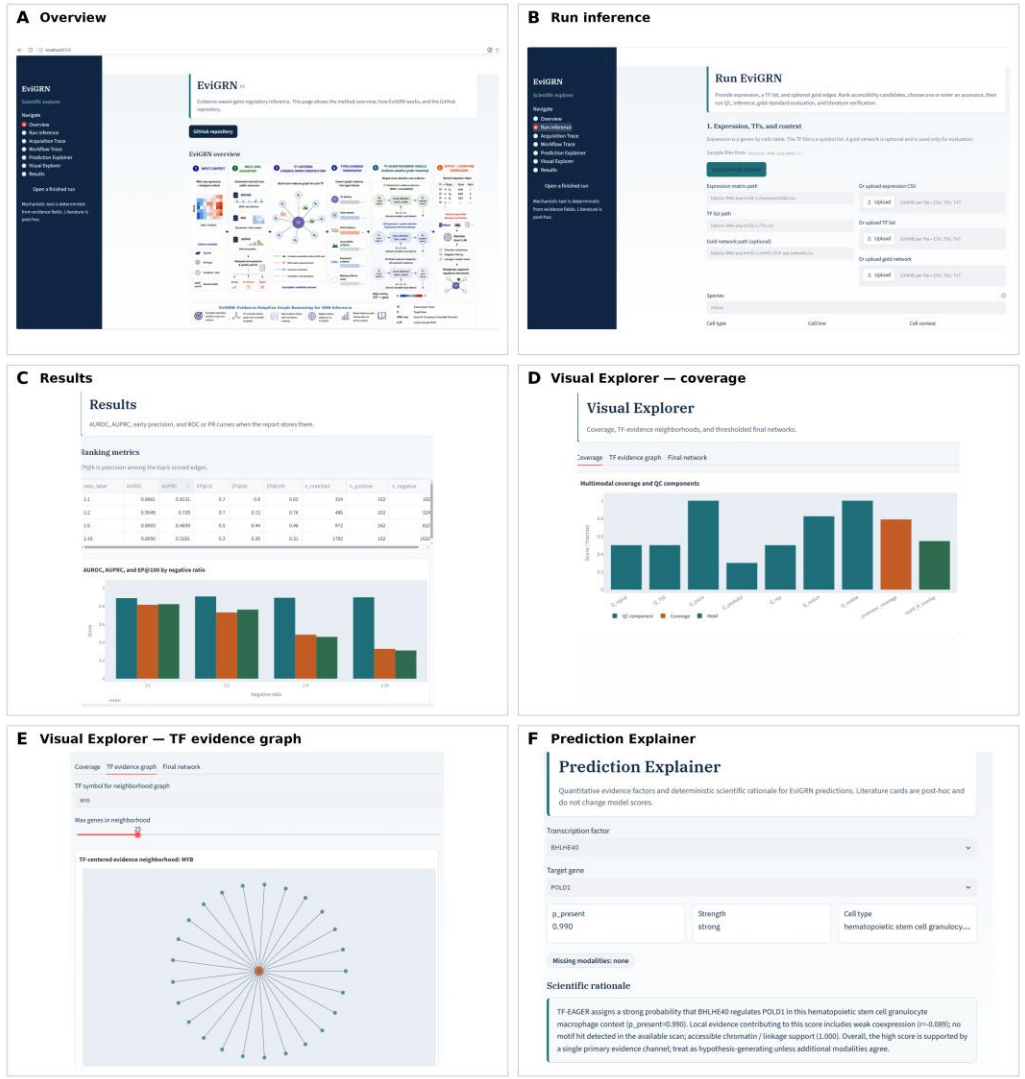

Acknowledgment: NSF awards 2343612 and 2525780, the University of Missouri, for its computational resources, and anonymous reviewers for all constructive feedback.