Beyond-Accuracy (Sparsed-) coVaraince Neural Network Recommender Systems

Author: Ivan Val. Bozhanin[1] | Responsible Professor: Elvin Isufi[1] | Supervisors: Andrea Cavallo[1], Chengen Liu[1]

[1]Delft University of Technology

The Big Picture: Why Better Recommendations?

Accuracy-optimised recommenders often pin users inside "filter-bubbles", limiting discovery and eventually causing content fatigue. We ask whether **coVariance Neural Networks (VNNs)** and their sparsified cousins (S-VNNs) can raise *Novelty* and *Diversity* while holding classical RMSE almost constant. **Research Questions:**

1. How do VNNs/S-VNNs perform on novelty & diverisity using covariance vs. precision matrices?
2. What's the impact of different sparsification techniques?
3. Is there a trade-off between rating prediction accuracy and beyond-accuracy metrics?

Dataset & Key Metrics

Statistic	Value
Users (U)	943
Movies (M)	1 682
Ratings	100 000
Split	80/10/10 (train/val/test)

RMSE: Average squared gap between the predicted rating and the true user rating. **Lower RMSE →more accurate predictions.**

Novelty: Mean *rarity score* of the items in a recommendation list, where rarity is "1 − (popularity of the item across the catalogue)". **Rewards surfacing lesser-known or newly discovered items; *higher* is better.**

Diversity: Average pairwise *hybrid dissimilarity* within the top-50, computed as 50% rating-cosine distance + 50% genre distance. **Rewards lists that span a broader range of themes, styles, or genres; *higher* is better**

Introducing VNNs & S-VNNs

- VNN (coVariance Neural Network):
  - Treats user/item relationships as covariance (precision) matrix.
  - Applies graph convolutions to learn patterns.
- S-VNN (Sparse VNN):
  - Sparsifies matrices, keeping only the important, meaningful correlations.
  - Reduces noice, improve efficiency, and potentially boosts novelty and diversity by focusing graph learning.

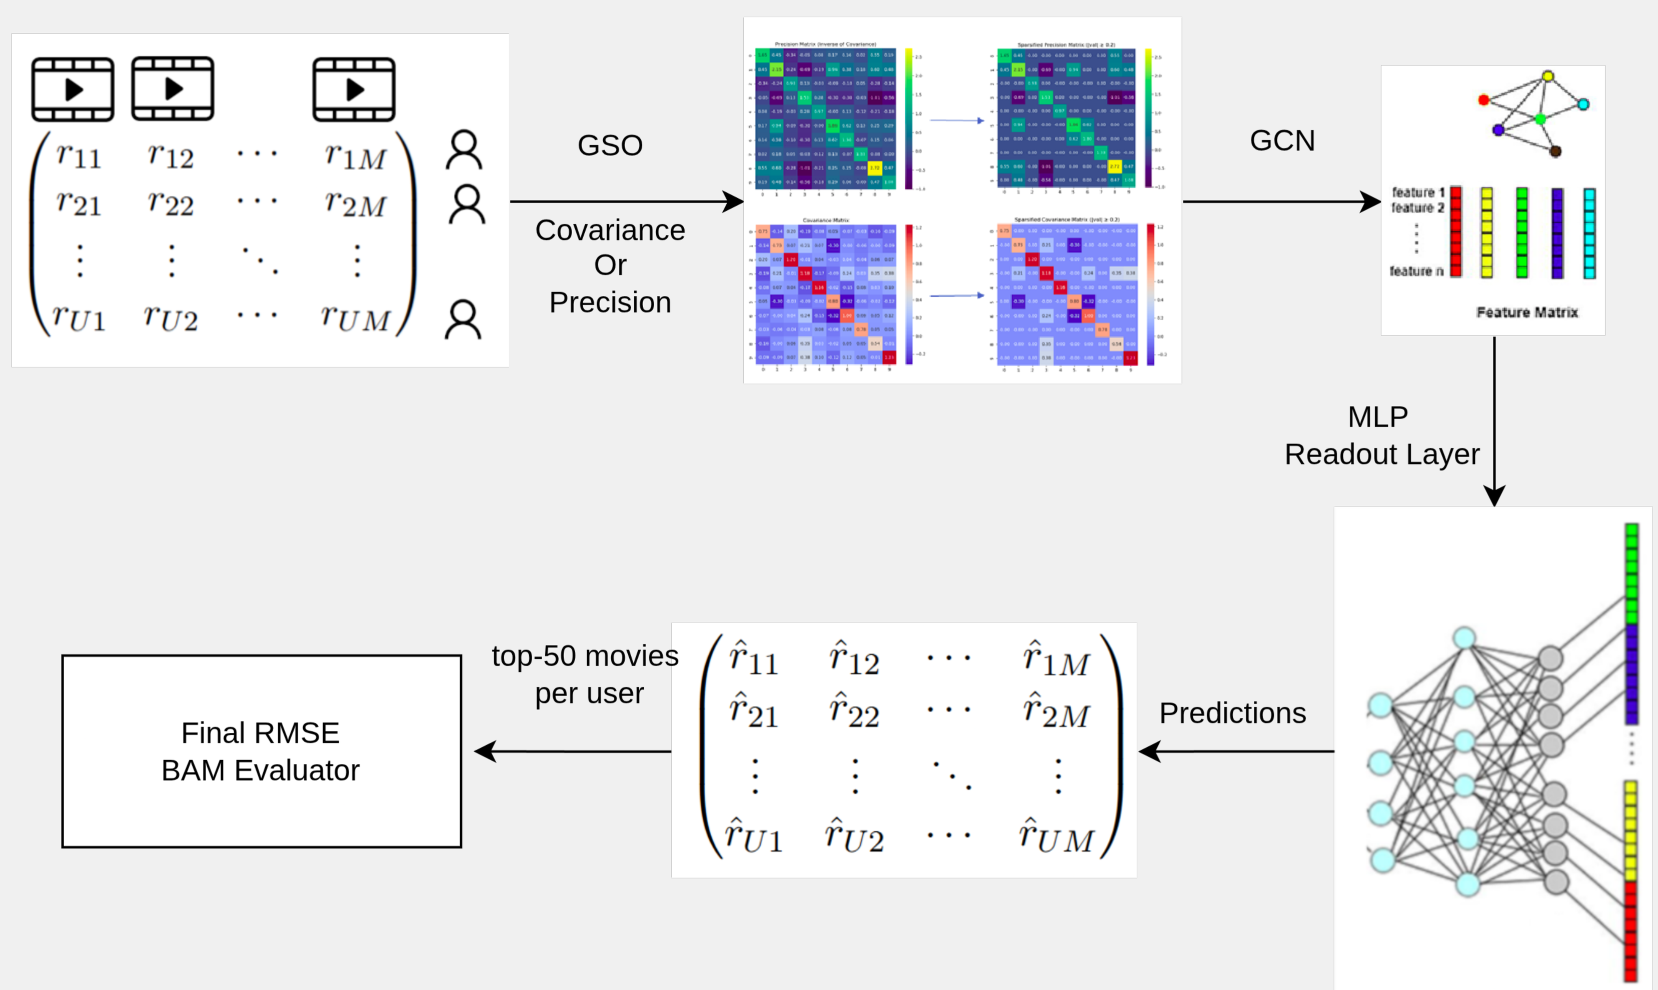

Outline of the model: Data → Graph Shift Operator (Cov/Prec, sparse/dense) → **SelectionGNN** filter bank ($K = 2$) → read-out → Top-50 ranking → BAM evaluation.

Accuracy vs. Beyond-Accuracy

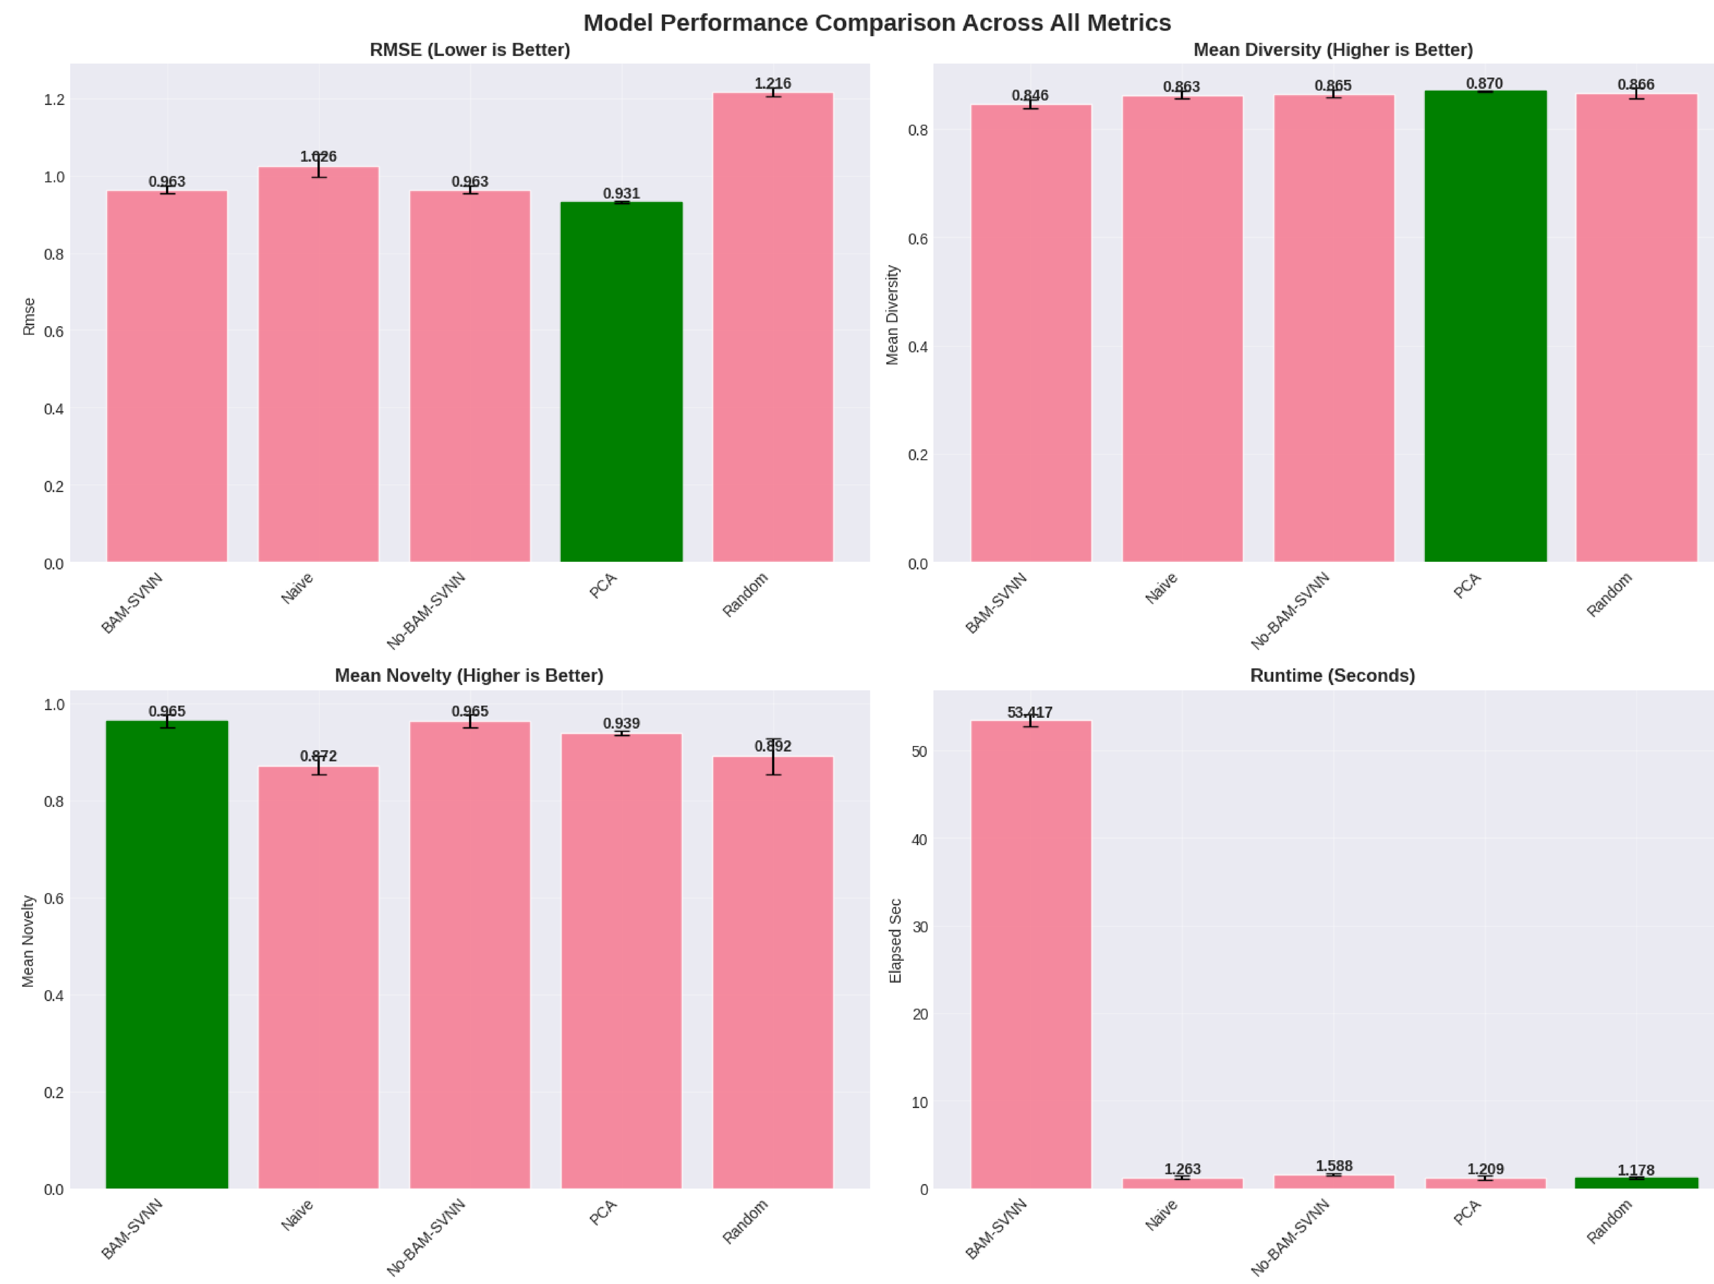

Mean ± 1 SE for five model families. Green bar = best per metric.

Trade-off Frontier

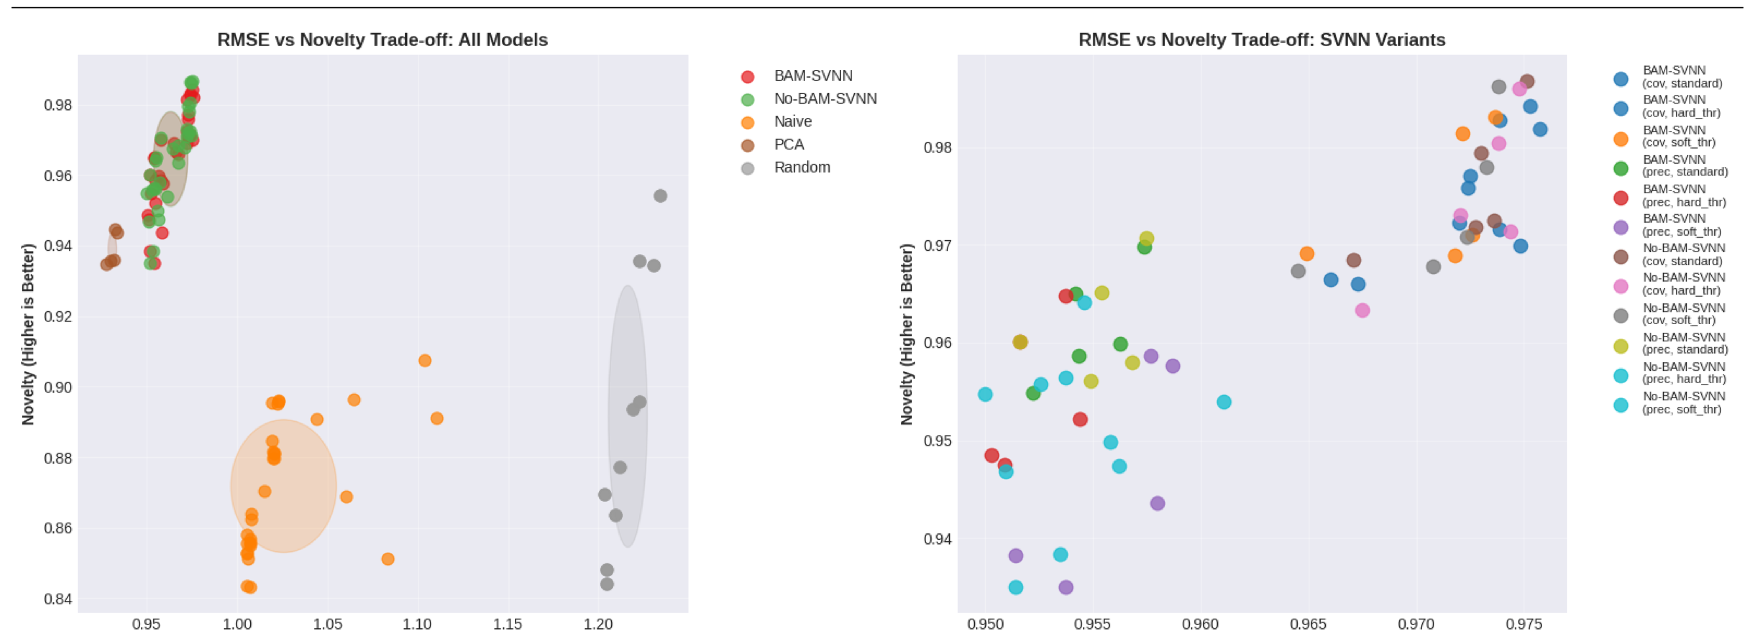

SVNN variants sit on the Pareto frontier: +2.8 pp Novelty for +0.03 RMSE versus PCA.

Key Observations

O1 Novelty Gain. Both SVNN flavours beat PCA by **+2.8 pp Novelty** while retaining near-parity RMSE.
O2 Graph Matters. Hard-thresholded *precision* graphs achieve the lowest SVNN RMSE (0.952); dense *covariance* graphs maximise Diversity (0.868).

Acknowledgments: Special thanks to my responsible professor - Elvin Isufi - and my supervisors - Andrea Cavallo, Chengen Liu.

Efficiency Matters

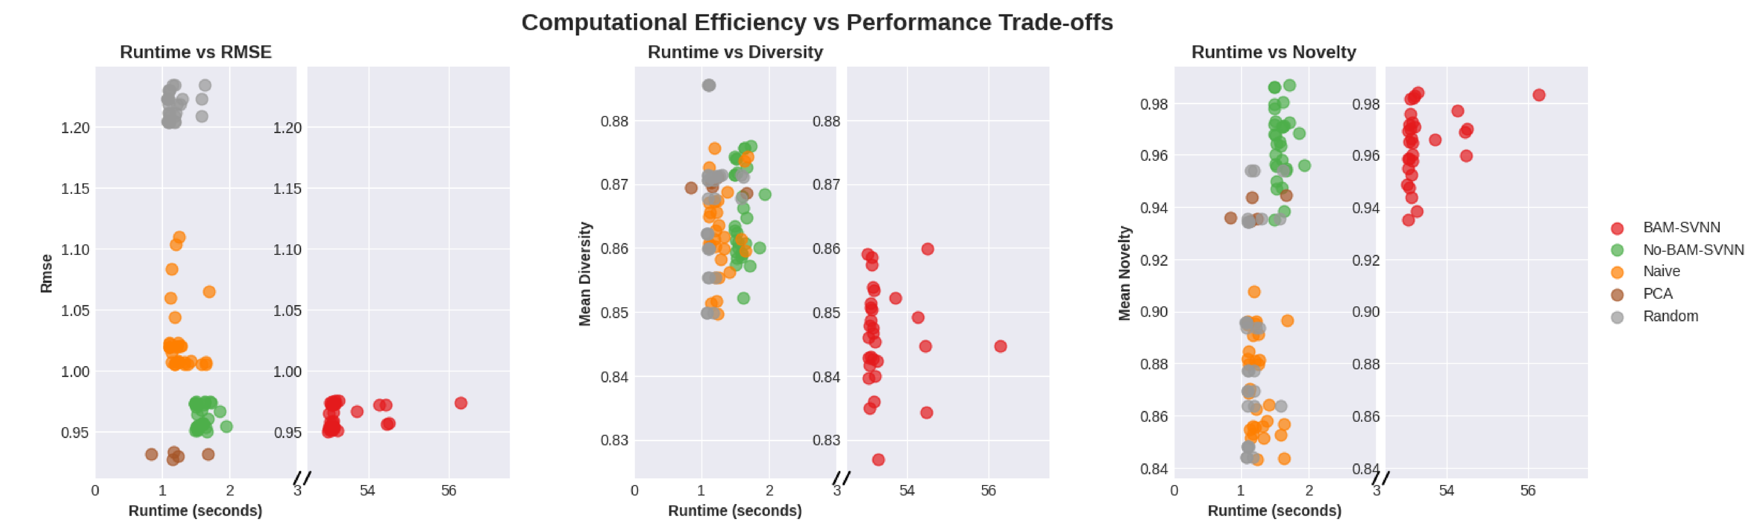

BAM-weighted loss multiplies runtime by ×33 yet shows no significant gain over RMSE-only SVNN.

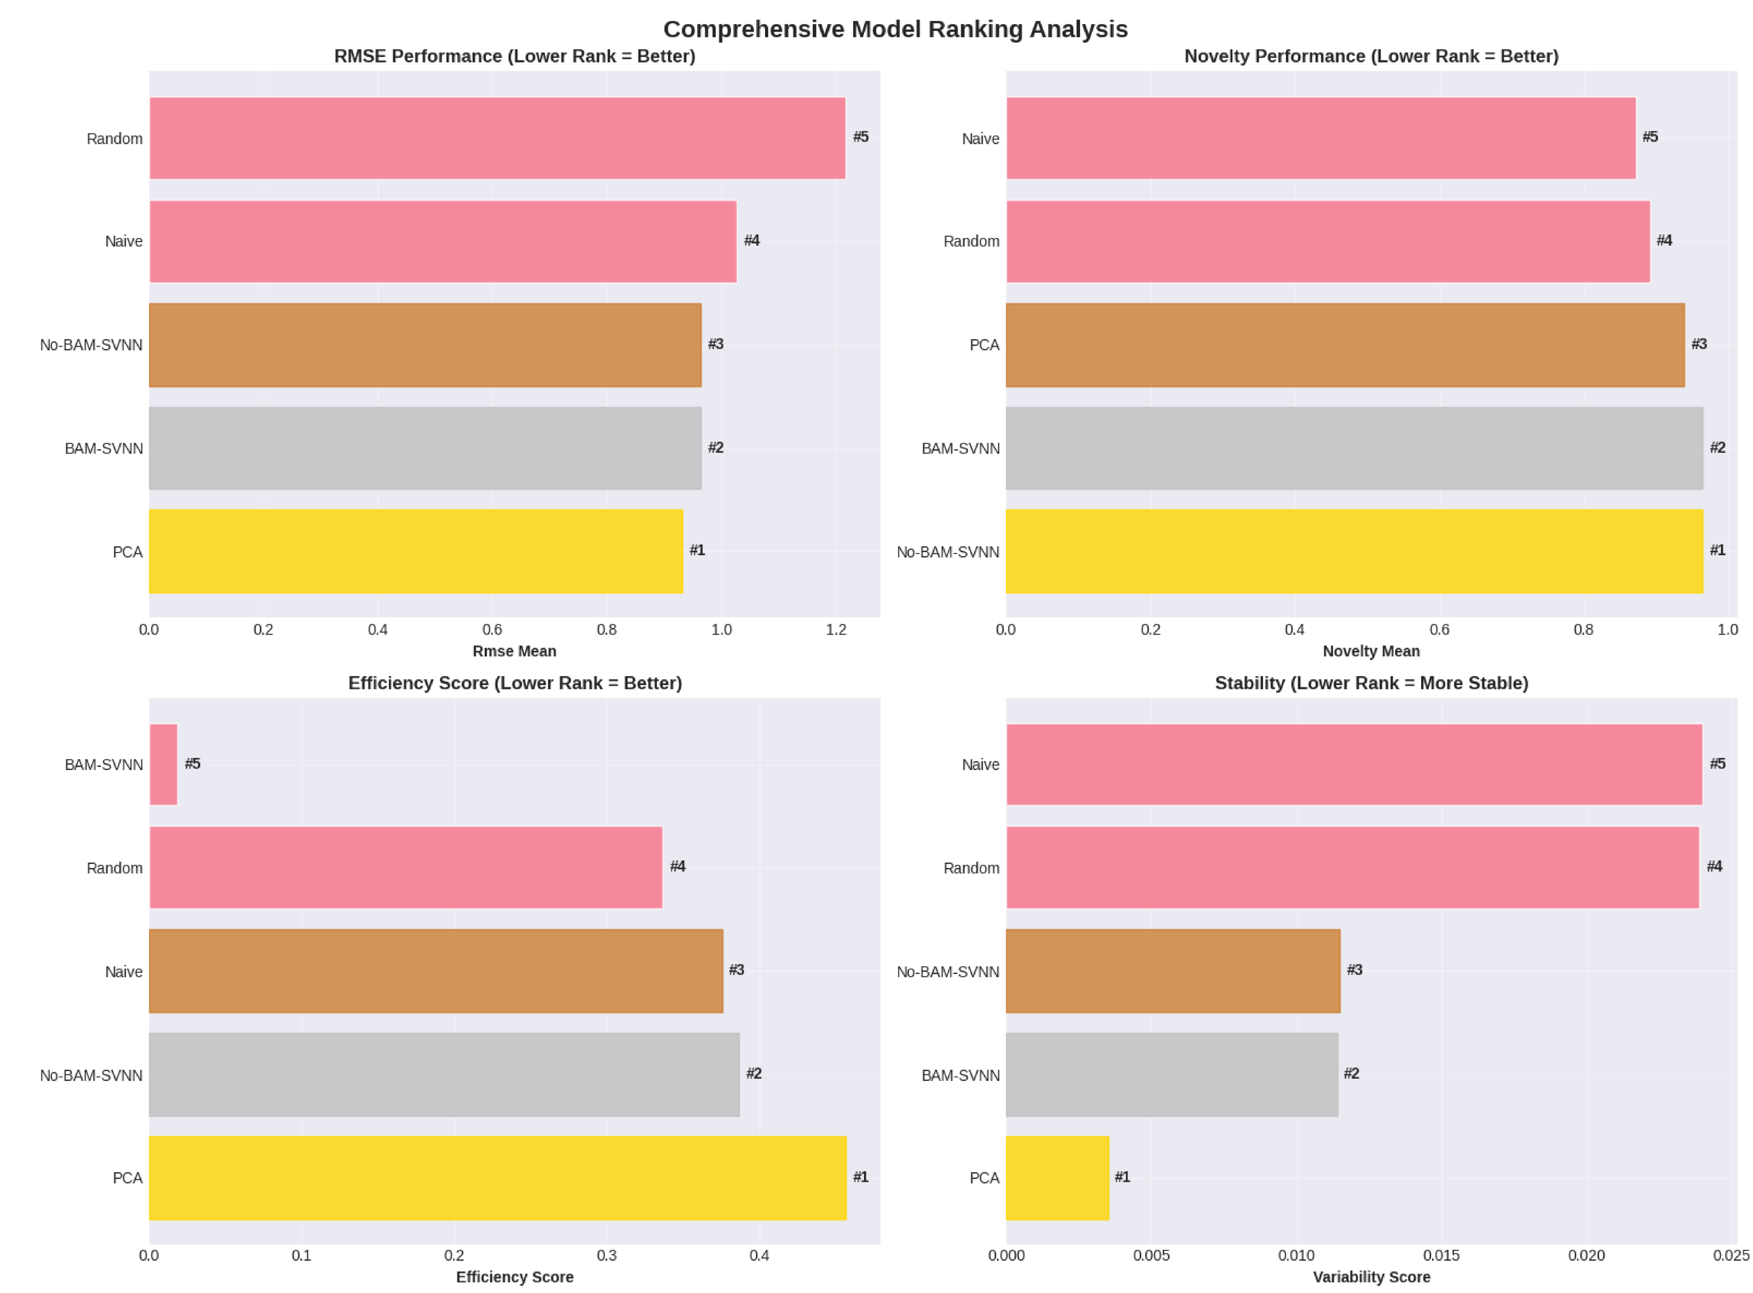

Future Work

Future Work

1. Validate on ML-1M, Amazon, YouTube implicit logs.
2. Optimise α and λ with multi-objective methods.
3. Learnable GSOs adaptive sparsification.
4. Deeper SVNN stacks; rank-based loss.
5. Online A/B test for perceived serendipity.

References & Acknowledgments

[1] Andrea Cavallo, Zhan Gao, and Elvin Isufi. Sparse Covariance Neural Networks, 2024.
[2] Saurabh Sihag, Gonzalo Mateos, Corey McMillan, and Alejandro Ribeiro. coVariance Neural Networks, 2023.
[3] Saúl Vargas and Pablo Castells. Rank and relevance in novelty and diversity metrics for recommender systems. 2011.