

Machine Learning Prediction of Molecular Initiating Events using Chemical Target Annotations and Gene Expression

J. L. Bundy¹, R. Judson¹, A.J. Williams¹, C. Grulke¹, I. Shah¹, L. J. Everett¹

¹ US EPA, Research Triangle Park, NC

The views expressed in this presentation are those of the author(s) and do not necessarily represent the views or policies of the Agency.



Logan J. Everett | everett.logan@epa.gov

Introduction

U.S. EPA is utilizing high-throughput transcriptomics (HTT) to screen chemicals for potential toxicity. There is an unmet need for methods to predict molecular initiating events (MIEs) that lead to toxic effects based on the transcriptional response to exposure, such as the primary protein targets of a chemical. To address this challenge, we are developing a machine learning based method that integrates HTT data and existing chemical-MIE labels to predict MIEs for unannotated chemicals.

Key points:

- Integrated LINCS L1000 CMAP gene expression compendium [1]
- Used RefChemDB database of chemical-protein target interactions [2]
- Trained binary classifiers on integrated data sets with the following parameters:
 - 52 MIE prediction models for MCF7-derived LINCS L1000 expression profiles
 - 3 Training Feature Types x 6 Classification Algorithms

Classifier Training Overview

The prediction of chemical bioactivity at the level of MIEs required the integration of existing chemical-MIE annotations and a large gene expression compendium (Figure 1).

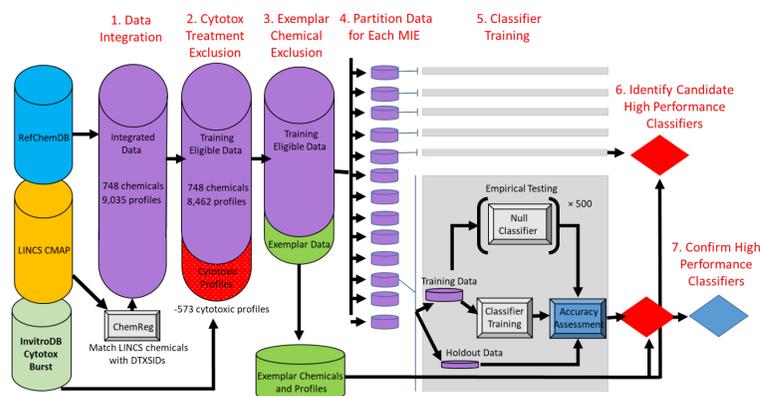


Figure 1. Data processing and classifier training workflow.

Method:

1. Chemical treatments associated with LINCS L1000 profiles were matched to EPA substance identifiers (DTXIDs) using ChemReg [3]
2. LINCS profiles corresponding to chemical concentrations above the "cytotoxic burst" value (InVitroDB) were excluded from the analysis
3. Exemplar chemicals associated with MIEs of interest were excluded from classifier training for downstream validation
4. Remaining data was partitioned into MIE-specific training data sets (see Figure 2)
5. Binary classifiers were trained independently for each of 52 distinct MIEs using the R package *caret*, 500 null classifiers (see Figure 4) were trained for each combination of MIE and classification algorithm
6. Classifiers that passed empirical significance testing were flagged as candidate high performance classifiers
7. Candidate high performance classifiers were validated with training excluded exemplar chemicals (see Figure 6)

Selection of Training Data

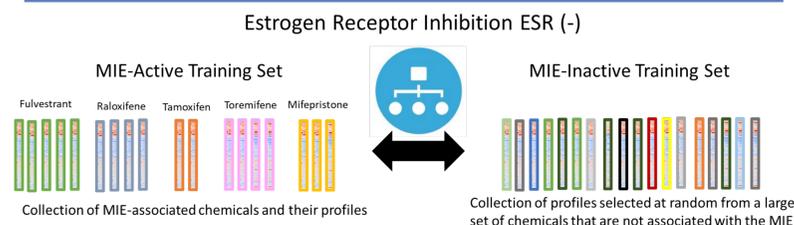


Figure 2. Example of training data structure for Estrogen Receptor inhibition. Binary classifiers were trained for each MIE using size-matched collections of LINCS L1000 gene expression profiles (represented by vertical bars) partitioned into a MIE-Active and MIE-Inactive category. MIE-Active profiles were associated with a chemical treatment that is linked to a given MIE in RefChemDB. MIE-Inactive profiles are selected at random from a collection of chemicals with no association with the given MIE in RefChemDB.

Classifier Optimization

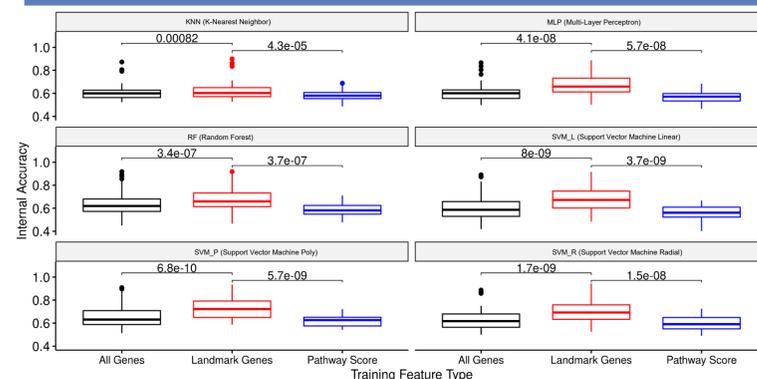


Figure 3. Comparison of internal accuracy distributions for classifiers trained each combination of training algorithm and training feature type. P-values are from a two-tailed paired Wilcoxon test.

- To optimize classifiers, we evaluated model performance across all the 3 types of gene expression feature sets and 6 classification algorithms (Figure 3).
- Classifiers were trained using 6 different algorithms
- Classifiers were trained using three different sets of features:
 1. All Genes:
 - ~12,000 transcripts, including both measured and inferred expression
 2. Landmark genes
 - ~1,000 transcripts that are directly measured in the L1000 assay
 3. Pathway scores
 - Canonical pathways from MSigDB [4] scored with gene set enrichment analysis [5], calculated from the "All Genes" features
- Cross-fold validation accuracies were compared for the 52 MIE classifiers trained on different feature types using a paired Wilcoxon test
- Landmark Gene based classifiers consistently out-performed "All Gene" and "Pathway Score" based classifiers, regardless of algorithm

Empirical Significance Testing

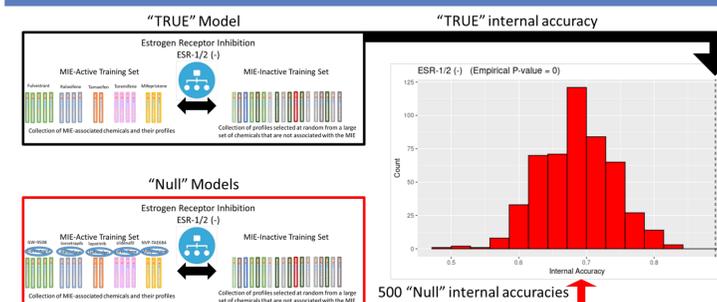


Figure 4. Example of empirical significance calculation using null model accuracy.

- Identified candidate high performance classifiers using an empirical significance testing approach
- Classifiers that generated an internal accuracy that is higher than 95% of their "null" counterparts (p-value < 0.05) were retained for further analysis
- 47 candidate high performance classifiers spanning 12 MIEs passed empirical significance testing and were then validated on exemplar chemicals

High-Performance Classifier Selection

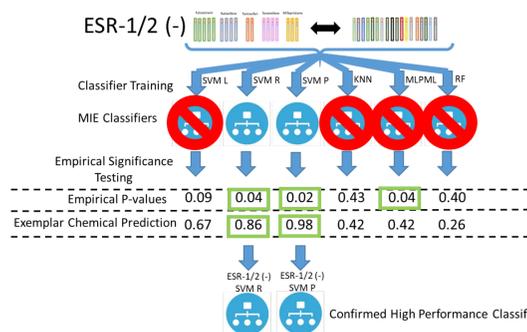


Figure 5. Schematic showing selection of high-performance classifiers. Classifiers were generated for the same MIE using 6 training algorithms. Classifiers were then filtered based on Empirical Significance testing (pvalue < 0.05) and exemplar chemical prediction (prediction > 0.75). Classifiers that met both criteria were termed "confirmed high performance classifiers".

- MIE activation predictions were generated for all MCF7-derived gene expression profiles in the LINCS L1000 CMAP data set
- 45 classifiers were validated using training-excluded exemplar chemicals. Retained classifiers must generate a prediction for their training-excluded exemplar chemical that is greater than 75% of the chemicals in the LINCS L1000 CMAP data set
- Predictions for high performance classifiers associated with the same MIE were averaged to generate ensemble predictions
- Confirmed high performance classifiers spanned 11 MIEs

References

1. Subramanian A, Narayan R, Corsello SM, Peck DD, Natoli TE, Lu XD, et al. A Next Generation Connectivity Map: L1000 Platform and the First 1,000,000 Profiles. *Cell*. 2017;171(6):1427.
2. Judson RS, Thomas RS, Baker N, Simha A, Howey XM, Marable C, et al. Workflow for Defining Reference Chemicals for Assessing Performance of In Vitro Assays. *Altox-Altern Anim Ex*. 2019;36(2):261-76.
3. Grulke CM, Williams AJ, Thillanadarajah I, Richard AM. EPA's DSSTox database: History of development of a curated chemistry resource supporting computational toxicology research. *Computational Toxicology*. 2019;12:100096.
4. Liberzon A, Subramanian A, Pinchback R, Thorvaldsdottir H, Tamayo P, Mesirov JP. Molecular signatures database (MSigDB) 3.0. *Bioinformatics*. 2011;27(12):1739-40.
5. Barbie DA, Tamayo P, Boehm JS, Kim SY, Moody SE, Dunn IF, et al. Systematic RNA interference reveals that oncogenic KRAS-driven cancers require TBK1. *Nature*. 2009;462(7269):108-U22.

Validated Classifier Predictions

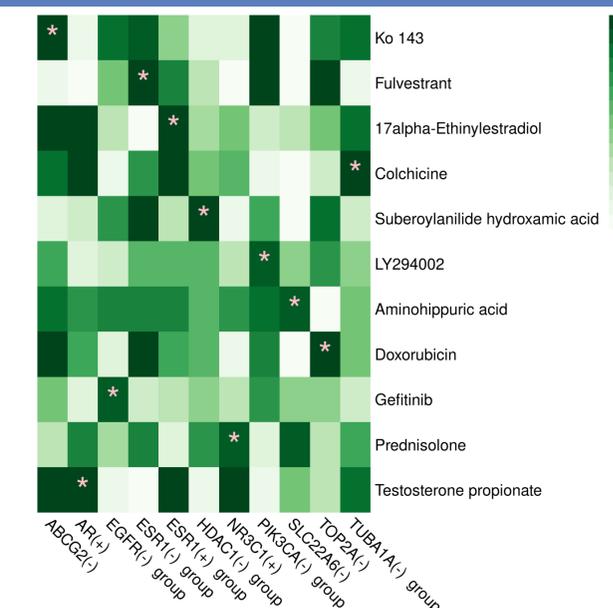


Figure 6. Heatmap showing MIE activation predictions from confirmed high-performance classifiers. Predictions shown for the 11 training-excluded exemplar chemicals (rows) from the 11 confirmed high-performance classifiers (columns). Shading indicates percent rank of each chemical's prediction relative to all other predictions for the same MIE. Dark green indicates affirmative prediction of MIE activation. * Indicates chemical-MIE linkage in RefChemDB.

- Confirmed high performance classifiers correct predicted MIE activation for their corresponding training-excluded exemplar chemicals
- Some classifiers predicted MIE activation for chemicals not annotated for the MIE
 - Possibly false positive predictions
 - Possibly the result of convergence of modeled MIEs onto key events activated by training-excluded exemplar chemicals

Discussion / Conclusions

- Integrated RefChemDB chemical-MIE annotations with LINCS chemical identifiers and gene expression profiles in a machine learning framework
- Trained binary classifiers to predict activation of 52 distinct MIEs
- Classifiers trained on landmark genes yielded the highest internal accuracy
- 47 classifiers that modeled MIEs significantly better than null models, 45 of which were validated with training-excluded exemplar chemicals
- 11 MIEs modeled with the remaining classifiers showed correctly predicted MIE activation of training-excluded exemplar chemicals
 - Some exemplar chemicals predicted as MIE active in the absence of literature annotations
- Findings suggest that ML-based methods for predicting MIEs may be helpful in prioritizing chemicals for further study based on transcriptomic profiling and may inform decisions on suitable cell-types for further screening