



Published in final edited form as:

*Pac Symp Biocomput.* 2021 ; 26: 273–284.

## Cross-modal representation alignment of molecular structure and perturbation-induced transcriptional profiles

Samuel G. Finlayson<sup>1,2,\*</sup>, Matthew B.A. McDermott<sup>2,\*</sup>, Alex V. Pickering<sup>3</sup>, Scott L. Lipnick<sup>3</sup>, Isaac S. Kohane<sup>3,†</sup>

<sup>1</sup>Department of Systems, Synthetic, and Quantitative Biology, Harvard Medical School, Boston, MA

<sup>2</sup>Department of EECS, Massachusetts Institute of Technology, Cambridge, MA

<sup>3</sup>Department of Biomedical Informatics, Harvard Medical School, Boston, MA

### Abstract

Modeling the relationship between chemical structure and molecular activity is a key goal in drug development. Many benchmark tasks have been proposed for molecular property prediction, but these tasks are generally aimed at specific, isolated biomedical properties. In this work, we propose a new cross-modal small molecule retrieval task, designed to force a model to learn to associate the structure of a small molecule with the transcriptional change it induces. We develop this task formally as multi-view alignment problem, and present a coordinated deep learning approach that jointly optimizes representations of both chemical structure and perturbational gene expression profiles. We benchmark our results against oracle models and principled baselines, and find that cell line variability markedly influences performance in this domain. Our work establishes the feasibility of this new task, elucidates the limitations of current data and systems, and may serve to catalyze future research in small molecule representation learning.

### Keywords

Representation Learning; Therapeutics; Gene Expression; Deep Learning; Information Retrieval

## 1. Introduction

Identifying molecules that are likely to have a specific biological effect is a cornerstone of drug discovery and a key component of efforts to achieve precision medicine. Classically, computational profiling of small molecules has centered on predicting affinities for specific biological targets, using tools ranging from biophysics-driven techniques such as molecular docking<sup>1</sup> to literature-mined annotations.<sup>2</sup> Small molecule modeling has also recently become a major area of interest in deep learning, a trend catalyzed by graph neural networks<sup>3</sup> and benchmarking datasets.<sup>4</sup> Graph neural networks allow for end-to-end

Open Access chapter published by World Scientific Publishing Company and distributed under the terms of the Creative Commons Attribution Non-Commercial (CC BY-NC) 4.0 License.

<sup>†</sup> isaac\_kohane@harvard.edu .

\* Co-first author

modeling of molecular graphs,<sup>5–8</sup> and have yielded state-of-the-art performance on certain tasks.<sup>9,10</sup> In addition, deep learning approaches have been used at a more global scale modeling cross-molecule relationships.<sup>11</sup> To date, deep learning efforts in this space have generally focused on two extremes: highly local, biochemical prediction problems, which test the model's ability to predict specific chemical properties, and more global, clinical modeling tasks, such as indication or side effect prediction. Missing from the field, however, are benchmark tasks between these two extremes, that test the ability of deep models to encode rich, general representations of a molecule's broad-spectrum effect on cellular biology.

In parallel to these developments, connectivity mapping has emerged as an alternative approach for drug development.<sup>12</sup> In connectivity mapping, compounds are foremost characterized not by individual chemical properties or downstream targets, but by the broad transcriptional effects they induce in cells. Connectivity mapping begins by first developing a large dataset of *perturbational signatures* of molecules by physically treating cell lines with these molecules, then measuring the resultant changes in gene expression. These datasets are then compared to one or more *query signatures*, which are typically differential gene expression (GE) signatures representing disease states that investigators hope to reverse. Various public datasets have been curated to enable these efforts,<sup>13,14</sup> and researchers have sought to use these for drug repurposing, precision medicine, and analysis of gene expression data in general.<sup>15–22</sup>

Connectivity mapping is promising because it can be used to search for new indications of drugs without making any specific *a priori* assumptions about their mechanism of action. However, the typical framework for connectivity mapping is limited by the fact that it can only query against drugs that have already been profiled using the transcriptional assay. In other words, connectivity mapping is – in principle – very flexible with respect to the disease signatures they accept as a query, but is transductive rather than inductive with respect to the target small molecule signatures. This is the perfect complement to structure-based computational chemistry, which is typically inductive to new drug structures but can only make predictions for diseases with known targets.

In this work, we combine these two fields, by using deep chemical embedders to learn the transcriptional space encoded by CMAP profiling. More specifically, we train coordinated networks to jointly embed chemical structures and perturbational gene expression profiles such that learned chemical representations are most similar to the encodings of the transcriptional patterns they induce.<sup>a</sup> Note that this task naturally fills the gap in inductive molecular modelling identified previously; by tasking the model to produce highly similar embeddings for chemical structures and the perturbational profiles they induce, we force the model to learn a *transcriptome-wide* reflection of the drug's action on the cell. We then evaluate these chemical representations by using gene expression signatures as queries into the embedding space and recovering their corresponding compounds. (See Figure 1). Crucially, the evaluation is set up such that the validation and test set compounds and cell

---

<sup>a</sup>Code available here: [https://github.com/sgfin/molecule\\_ge\\_coordinated\\_embeddings](https://github.com/sgfin/molecule_ge_coordinated_embeddings)

lines are not used in training, which allows us to test the ability of the model to generalize to new drugs and cell lines.

In the rest of this work, we first offer some background on joint embedding alignment, then detail the methods used in this work. Finally, we walk through the results and discussion of these experiments, then close with concluding thoughts. A version of this work, with supplementary material present, can be found here: [https://github.com/sgfin/molecule\\_ge\\_coordinated\\_embeddings/blob/master/paper\\_08\\_2020.pdf](https://github.com/sgfin/molecule_ge_coordinated_embeddings/blob/master/paper_08_2020.pdf).

## 2. Background

In *multi-view representation alignment*, embeddings of two associated data modalities are learned separately but in a coordinated manner, such that the resulting embeddings are similar. These methods have been used in comparing images to text but also in other domains.<sup>23,24</sup> In this work, we learn aligned representations such that small molecules are embedded in close proximity to the differential gene expressions they induce. Multi-view representation alignment can be achieved through a variety of methods, including classical methods, such as *canonical correlation analysis* (CCA)<sup>25</sup> and methods using distance, similarity, correlation, or ranking based penalties during training.<sup>26</sup> Ranking-based methods for multi-view representation alignment, such as that described by Deng et al,<sup>27</sup> allow the incorporation of ranking information into the training procedure, which may be important in tasks such as gene expression where perturbation signals may be small relative to baseline state. In addition, the field of rank-based embedding learning is intertwined with a broader literature of uni-modal embedding learning, which pioneered such architectures as Twin<sup>28,29</sup> and Triplet networks,<sup>30</sup> which optimize embeddings to bring similar data together while driving dissimilar data apart. An analysis of best practices of these architecture can be found in Wu et al.<sup>31</sup>

## 3. Methods

### 3.1. Dataset & Tasks

**Data Acquisition and Subsetting**—All data in this study comes from the LINCS Consortium/NIH Next-Generation Connectivity Map Level 3 L1000 data.<sup>14</sup> This dataset features 978-dimensional gene expression profiles from a variety of human cell lines treated with chemical and genetic perturbations. To ensure support over possible drugs, our data cut uses the most frequent 8 cell lines split into train, validation, and test sets such that no cell line or drug in the train set appears in the validation or test sets. To mitigate non-random missingness, we included only drugs assayed in all cell lines, and limited experiments to those incubated with small molecules for 24 hours at a dose of 10 $\mu$ m. Final statistics of these data are shown in Table A1. For drug structures, we used the SMILES<sup>32</sup> structures provided by LINCS, canonicalized using RDKit.<sup>33</sup>

**Preprocessing and Feature Engineering**—Gene expression intensity values from the training, validation, and test sets were centered and scaled at the gene-level based on the mean and standard deviation of each gene intensity across the training set.

We augmented each gene expression profile with three additional sets of features: the

corresponding gene expression intensities from a control signature on the same plate, the  $\log_2$  fold-change between the perturbation and control signatures, and the difference between these gene expression signatures. For use in our baseline and oracle models, we also computed numerical representations of each small molecule: Morgan extended-connectivity fingerprints<sup>34</sup> and the output of the ChemProp network.<sup>6</sup>

**Detailed Task Description**—Our goal is to learn embedders which map molecular structures and gene expression profiles into a vector space such molecular structure embeddings are close to the gene expression profile they induce while being far from other gene expression profiles (Figure 1). Formally, given a collection of gene expression signatures  $\mathcal{G}$ , chemical structures  $\mathcal{M}$ , and similarity function  $\text{Sim} : \mathbb{R}^d \times \mathbb{R}^d \rightarrow \mathbb{R}$ , we seek to learn a gene expression embedder  $E_g : \mathcal{G} \rightarrow \mathbb{R}^d$  and chemical embedder  $E_m : \mathcal{M} \rightarrow \mathbb{R}^d$  to maximize  $\text{Sim}(E_g(g_i); E_m(m_j))$  while simultaneously minimizing  $\text{Sim}(E_g(g_i), E_m(m_{-j}))$ , where gene expression  $g_i$  was induced by molecule  $m_j$ . Unless otherwise specified, the similarity function can be assumed to be Pearson Correlation in our experiments. Across our baseline and oracle methods, we realize many variants of  $E_g$  and  $E_m$ .

### 3.2. Baseline and Oracle Methods

**Nearest Neighbor Baseline**—Nearest-neighbor (NN) methods have been previously shown to establish strong baselines for machine learning tasks on the L1000 data.<sup>35,36</sup> In our cross-modal, information retrieval (IR) context, traditional NN methods are not applicable, so we employ the following “double NN” baseline: given a gene expression profile as a query, we first identify the nearest gene expression profile in the train set and look up its corresponding small molecule. We then take this small molecule (from the train set) as a query, and return the most structurally similar drug from the *test* set as our final prediction.

In particular, given a mapping  $G2M : \mathcal{G} \rightarrow \mathcal{M}$  from gene expression profiles to the small molecule that induced them, and a molecular embedding  $E_m$  (which may include molecular fingerprints, Chemprop embeddings, or embeddings learned from other models), we define embedder  $E_g : g_{\text{query}} \mapsto E_m(G2M(\arg\max_{g_{\text{tr}} \in \mathcal{G}_{\text{train}}} \text{Sim}(g_{\text{query}}, g_{\text{tr}})))$ . Then, we perform information retrieval (IR) analyses with such embedders as usual.

**Canonical Correlation Analysis Baseline**—Given training matrices of transcriptional  $\mathcal{G}_{\text{train}}$  and molecular  $\mathcal{M}_{\text{train}}$  encodings, we can learn a set of linear mappings  $E_g : \mathcal{G}_{\text{train}} \rightarrow \mathbb{R}^d$  and  $E_m : \mathcal{M}_{\text{train}} \rightarrow \mathbb{R}^d$  via  $d$ -dimensional CCA such that these mappings optimize the correlation between elements of  $\mathcal{G}_{\text{train}}$  and  $\mathcal{M}_{\text{train}}$ .

Note that this procedure requires a default numerical representation for molecules, which, as with other methods, can be either fingerprints, ChemProp embeddings, or learned embeddings by our learning model (described below). CCA can also be performed atop other embedding systems to further optimize embedding results. CCA was performed using SciKit Learn,<sup>37</sup> using 50 components, chosen to optimize validation set performance via a grid-search over a range of 5–125 components, run for 1000 iterations to ensure convergence.

**Oracle Models**—The central objective of our task is to learn small molecule embeddings that can stand in as surrogates for their corresponding gene expression signatures. To provide a rough upper-bound for expected performance on this task, we also implemented two “oracle” models, each of which queries test set GE signatures against pseudo-“chemical embeddings” that are in reality the average GE signatures from each test set drug when it was measured on either (1) the *train set cell lines*,<sup>b</sup> to simulate an embedder that perfectly associates all structures to perturbational profiles, but cannot generalize beyond the train set cell lines, or (2) the *test set cell lines*, which simulates a model of the same capabilities but able to generalize perfectly to the test set as well. These oracle models are still dependent on the underlying gene expression signature representation, so further innovation could offer improved upper bounds for this task. Formally, given  $G2M$  mapping gene expression profiles to their corresponding perturbing molecule, we define oracle embeddings  $E_g^{\text{train}}: g_{\text{query}} \mapsto \text{Avg}(\{g_i \in \mathcal{G}_{\text{train}} \mid G2M(g_i) = G2M(g_{\text{query}})\})$ , and  $E_g^{\text{test}}: g_{\text{query}} \mapsto \text{Avg}(\{g_i \in \mathcal{G}_{\text{test}} \mid G2M(g_i) = G2M(g_{\text{query}})\})$ .

### 3.3. Deep Coordinated Metric Learning Approach

For our learned model, we realize  $E_g$  as a self-normalizing neural network (of size dictated by hyperparameter search), and  $E_m$  as a directed message-passing neural network (D-MPNN), initialized by the Chemprop system, followed by a feed-forward output layer whose shape was dictated via hyperparameter search.<sup>6</sup> To train these architectures, we use a margin-based quadruplet loss, building on Wu et al’s adaptive margin loss.<sup>31</sup> The base of the adapted margin loss is defined over two data points  $i$  and  $j$  as  $\text{mar}_{\alpha,\beta} := (\alpha + y_{i,j}(D_{ij} - \beta))_+$ , where  $D$  is distance function (here euclidean distance),  $\alpha$  defines a permissible margin of separation,  $\beta$  controls the boundary between positive and negative pairs, and  $y_{i,j}$  is an indicator variable equal to 1 if  $i$  and  $j$  are of the same class and 0 otherwise ( $\alpha$  and  $\beta$  were tuned as hyperparameters).

Given two pairs of matching gene expression and molecular structure embeddings,  $(g_A, m_A)$ ,  $(g_B, m_B)$ , our quadruplet loss is defined as the sum of the margin losses between all cross-modality pairs of embeddings:  $\ell_{\text{quad}} = \text{mar}_{\alpha,\beta}(g_A, m_A) + \text{mar}_{\alpha,\beta}(g_A, m_B) + \text{mar}_{\alpha,\beta}(g_B, m_A) + \text{mar}_{\alpha,\beta}(g_B, m_B)$ . The network is thus optimized to bring the positive embedding within the margin of the anchor and negative embedding outside the margin. For sampling these two pairs (an analog of negative sampling for a more traditional triplet network), we first sample one matching pair, choose the molecular structure for the other pair based on the distance-weighted negative sampling scheme described in Wu et al, which was successful with their margin-based approach,<sup>31</sup> then fill in the other gene expression profile to match the sampled molecular structure. To make this process computationally efficient we pre-computed the average distance in *average post-perturbational gene expression space* between every pair of small molecule structures in the dataset. We additionally tried other losses, including two varieties of traditional triplet losses, and a quintuplet loss, but ultimately found the quadruplet loss to be most performant via our hyperparameter search.

<sup>b</sup>Note that we can do this as we limited our choice of drugs to those that were measured in all 8 cell lines, even though our actual data split prohibits training on any drug that appears in the validation or test sets.

**Training and Hyperparameter Selection**—Each model was trained on a Nvidia GeForce GTX 1080 GPU. Early stopping was used to select the model with the best mean reciprocal rank on the validation set. Hyperparameter tuning via the Bayesian Hyperopt library<sup>38</sup> was performed over a wide range of possible hyperparameters, including network depth and width (parametrized by the size of the first hidden layer and a growth rate), learning rate, number of epochs, batch size, margin and  $\beta$  parameters, triplet model orientation (i.e., gene first or compound first), activation function/network type (e.g., SNNN vs. unconstrained fully connected network), and dropout, with no early stopping for hyperparameter search runs. The optimal hyperparameters from this search are shown in Appendix Section 7.2.

### 3.4. Experiments

We designed a range of experiments with two purposes: First, we sought to evaluate if and how our deep coordinated representation learning method offers improvements over principled baselines. This entails a quantitative performance comparison against baseline methods and ablated versions of our model. Second, we introspect into the representations learned by training on this new task, to better understand the challenges and utility of the general framework. This entails a quantitative performance comparison against oracle models, a statistical analysis probing the ability of our models to generalize to new structures, and a qualitative exploration of the changes in chemical representations that are induced by our training scheme.

Quantitative performance analyses began by computing the embeddings of gene expression signatures and chemical structures in the test set, using the baseline, oracle, and deep coordinated methods defined in Sections 3.2 and 3.3. Using each embedded gene expression signature as a query, we ranked all chemical structures in the test set based on their proximity to that gene expression signature in the embedding space. These rankings were then used to compute standard information-retrieval metrics (precision-recall curves, MR, MRR, and Hits at/H@ 10 or 100). For our ablation analyses, we repeated the above experiments using various combinations of raw and learned GE and chemical representations.

In addition to the information retrieval analyses, we probed the generalizability of our representations by analysing the statistical relationship between the average retrieval performance for each chemical structure and the structural similarity to the most similar chemical in the training set. Similarity was measured via the Tanimoto distance between all pairs of molecular fingerprints in the train and test set. We further examined performance vs. chemical specificity, using the number of genes that a molecule, on average, affected as a measure of specificity. Finally, we visualized the latent space of our chemical embedder (versus their pre-trained representations learned by ChemProp), and noted the relative position of drugs with the same mechanism of action (MOA) in each latent space.

## 4. Results & Discussion

### 4.1. Quantitative IR Experiments

**Baseline and Proposed Method**—Information retrieval results from the baseline and proposed methods are reported in Figure 2. This figure shows that our model variants offer significant performance improvements over either of the CCA or NN baselines, and even approach the performance of the train-set only oracle model. All models fall dramatically short of the test-set generalizability oracle, which indicates that while our tasks offer significant improvement over baseline models here, there are still major gains possible, primarily by focusing on improving the generalizability to the novel cell-types of the test set.

In addition, we show various ablation studies over the baseline models in Table 1 to probe what gene or molecular representations would make them better or worse. Strikingly, we can note that uniformly using aligned representations (meaning representations based on our multi-view alignment neural network architecture) offers significant improvements over other representations, indicating that even with a baseline approach such as the double nearest neighbor (D-NN) model, improvements to the embedding quality translate to notable improvements to IR performance. Notably, this is true both for GE and Chemical embedders, with Aligned-Aligned representations yielding optimal performance for both D-NN and CCA query mechanisms. Additionally, it is also clear that CCA is the preferred query metric, over either raw correlation (Corr) based lookups or D-NN based lookups.

**Oracle Model Analysis**—Results from the oracle models are reported in Figure A1 and Table 2. As expected, oracle models using GE signatures from the test cell line greatly outperformed those using signatures from the train cell lines. This stark difference suggests that one of the largest barriers to performance here is the generalization gap between different cell lines. This further motivates for the curation of larger, cell-line heterogeneous datasets in the future.

In addition, aligned embeddings modestly improved the performance of test set oracles, and greatly improved the performance of train set oracles, consistent with the GE embedders learning slightly more generalizable representations. Of note, our proposed model achieved comparable results as the oracle model that leveraged raw GE signatures from the train cell lines. More specifically, our approach yielded slightly worse than the oracle on metrics (MRR, H@10) that emphasize early rankings, and slightly better on metrics (MR, H@100) that focus on more aggregate results. This is also apparent in the precision-recall plot, which shows the aligned embeddings curve starting out slightly below that of the raw/train oracle curve but then moving rapidly above it as further results are considered.

### 4.2. Introspection Analyses

Figure 3A contains the results of our experiments comparing performance to distance from the training set. Regardless of the measure of chemical similarity, compound retrieval performance was inversely correlated with distance from the training set. As can be seen in Supplementary Figure A2, the same trend held with learned gene expression embeddings, and was present but much weaker using raw gene expression profiles.

Figure 3B depicts the relationship between the transcriptional specificity of a compound and its ability to be retrieved using our analysis. As can be seen, there is a mild negative correlation, implying that molecules that affect the expression of many genes are easier to retrieve using this approach. Note that this observation is concordant with our findings on the difficulty of generalizing to new cell lines – drugs that affect a small, targeted set of genes are more likely to be cell line specific, and as our model is forced to surmount a significant generalization gap in evaluation, such cell-line specific signals are largely wiped out.

In addition, our analysis of the changes induced in the embedding space, shown in Supplementary Figure A4, reveal that our model's embeddings of molecules appear to better cluster shared MOAs than do the raw ChemProp embeddings, from which our model is initialized. This suggests that, as hypothesized regarding the nature of this task, our model is learning rich representations of the underlying molecules, though additional work remains to investigate this effect more thoroughly.

### 4.3. Future Work

We see several opportunities for further work on this task. First, expanding our data coverage, across molecules, cell lines, dosages, and treatment durations will allow us to measure and improve generalizability here. Second, exploring additional strategies to use the oversampled nature of these data (e.g., ensembling together control and perturbational signatures to reduce variance) could be beneficial. Third, a more robust exploration of model architectures, losses, and deep metric learning/negative sampling methods, could offer improvements here.<sup>26</sup> Additionally, other styles of multimodal embedding could be explored, such as the use of cycle generative adversarial network, which in particular would enable us to adopt a semi-supervised approach.<sup>39,40</sup> The use of interpretability methods, particularly those used for graph analyses,<sup>41</sup> as well as additional studies interrogating how our model's performance changes with the amount of available training data could also be insightful here. Fourth, we recommend exploring methods to improve cell-line generalizability, e.g. incorporating information across many cell lines when forming predictions. Finally, we also note that while our analyses only examine small-molecule therapeutics, similar methods could also be applied to other modalities, such as RNA-based therapies.

## 5. Conclusion

We present a new task: cross-modal multi-view alignment between drug structures and perturbational gene expression profiles, which links molecular structure to an objective, functional readout of drugs with very broad biomedical relevance. We profile state-of-the-art representation learning methods on this new task, and inspect the learned chemical embeddings. We find that this modeling task induces an embedding space reflective of drug mechanism of action – which is not explicitly included in the training regime – and see modest generalization to both new structures and a new biological environment. Our oracle experiments demonstrate major performance gaps when trying to generalize to new tissues.



We hope that this new benchmark task will catalyze future research and ultimately help enable a rapid, in silico compound prioritization methods.

## Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

## Acknowledgements

The authors thank Connor Coley and Kyle Swanson for providing the pretrained chemical embedder from Yang et al.<sup>6</sup> S.G.F. was supported by training grant T32GM007753 from the National Institute of General Medical Science. M.B.A.M. was funded in part by National Institutes of Health: National Institutes of Mental Health grant P50-MH106933 as well as a Mitacs Globalink Research Award. The content is solely the responsibility of the authors.

## References

1. Hansson T, Oostenbrink C and van Gunsteren W, Molecular dynamics simulations, Current opinion in structural biology 12, 190 (2002). [PubMed: 11959496]
2. Krallinger M, Rabal O, Lourenco A, Oyarzabal J and Valencia A, Information retrieval and text mining technologies for chemistry, Chemical reviews 117, 7673 (2017). [PubMed: 28475312]
3. Battaglia PW, Hamrick JB, Bapst V, Sanchez-Gonzalez A, Zambaldi V, Malinowski M, Tacchetti A, Raposo D, Santoro A, Faulkner Ret al., Relational inductive biases, deep learning, and graph networks, arXiv preprint arXiv:1806.01261 (2018).
4. Wu Z, Ramsundar B, Feinberg EN, Gomes J, Geniesse C, Pappu AS, Leswing K and Pande V, Moleculenet: a benchmark for molecular machine learning, Chemical science 9, 513 (2018). [PubMed: 29629118]
5. Kearnes S, McCloskey K, Berndl M, Pande V and Riley P, Molecular graph convolutions: moving beyond fingerprints, Journal of computer-aided molecular design 30, 595 (2016). [PubMed: 27558503]
6. Yang K, Swanson K, Jin W, Coley CW, Eiden P, Gao H, Guzman-Perez A, Hopper T, Kelley B, Mathea Met al., Analyzing learned molecular representations for property prediction, Journal of chemical information and modeling (2019).
7. Lo Y-C, Rensi SE, Tornig W and Altman RB, Machine learning in chemoinformatics and drug discovery, Drug discovery today 23, 1538 (2018). [PubMed: 29750902]
8. Ramsundar B, Kearnes S, Riley P, Webster D, Konerding D and Pande V, Massively multi-task networks for drug discovery, arXiv preprint arXiv:1502.02072 (2015).
9. Faber FA, Hutchison L, Huang B, Gilmer J, Schoenholz SS, Dahl GE, Vinyals O, Kearnes S, Riley PF and Von Lilienfeld OA, Prediction errors of molecular machine learning models lower than hybrid dft error, Journal of chemical theory and computation 13, 5255 (2017). [PubMed: 28926232]
10. Gilmer J, Schoenholz SS, Riley PF, Vinyals O and Dahl GE, Neural message passing for quantum chemistry, Proceedings of the 34th International Conference on Machine Learning-Volume 70, 1263 (2017).
11. Zitnik M, Agrawal M and Leskovec J, Modeling polypharmacy side effects with graph convolutional networks, Bioinformatics 34, i457 (2018). [PubMed: 29949996]
12. Musa A, Ghoraie LS, Zhang S-D, Glazko G, Yli-Harja O, Dehmer M, Haibe-Kains B and Emmert-Streib F, A review of connectivity map and computational approaches in pharmacogenomics, Briefings in bioinformatics 19, 506 (2017).
13. Lamb J, Crawford ED, Peck D, Modell JW, Blat IC, Wrobel MJ, Lerner J, Brunet J-P, Subramanian A, Ross KNet al., The connectivity map: using gene-expression signatures to connect small molecules, genes, and disease, science313, 1929 (2006). [PubMed: 17008526]
14. Subramanian A, Narayan R, Corsello SM, Peck DD, Natoli TE, Lu X, Gould J, Davis JF, Tubelli AA, Asiedu JKet al., A next generation connectivity map: L1000 platform and the first 1,000,000 profiles, Cell171, 1437 (2017). [PubMed: 29195078]

15. Musa A, Tripathi S, Kandhavelu M, Dehmer M and Emmert-Streib F, Harnessing the biological complexity of big data from lincs gene expression signatures, *PloS one* 13, p. e0201937 (2018). [PubMed: 30157183]
16. Liu T-P, Hsieh Y-Y, Chou C-J and Yang P-M, Systematic polypharmacology and drug repurposing via an integrated 11000-based connectivity map database mining, *Royal Society open science* 5, p. 181321 (2018). [PubMed: 30564416]
17. Clark NR, Hu KS, Feldmann AS, Kou Y, Chen EY, Duan Q and Maayan A, The characteristic direction: a geometrical approach to identify differentially expressed genes, *BMC bioinformatics* 15, p. 79 (2014). [PubMed: 24650281]
18. Donner Y, Kazmierczak S and Fortney K, Drug repurposing using deep embeddings of gene expression profiles, *Molecular pharmaceutics* 15, 4314 (2018). [PubMed: 30001141]
19. Dincer AB, Celik S, Hiranuma N and Lee S-I, Deepprofile: Deep learning of cancer molecular profiles for precision medicine, *bioRxiv*, p. 278739 (2018).
20. Rampásek L, Hidru D, Smirnov P, Haibe-Kains B and Goldenberg A, vae Dr.: improving drug response prediction via modeling of drug perturbation effects, *Bioinformatics* 35, 3743 (2019). [PubMed: 30850846]
21. Cheng J, Xie Q, Kumar V, Hurlle M, Freudenberg JM, Yang L and Agarwal P, Evaluation of analytical methods for connectivity map data (World Scientific, 2013) pp. 5–16.
22. Lopez R, Regier J, Cole MB, Jordan MI and Yosef N, Deep generative modeling for single-cell transcriptomics, *Nature methods* 15, 1053 (2018). [PubMed: 30504886]
23. Hsu T-MH, Weng W-H, Boag W, McDermott M and Szolovits P, Unsupervised multimodal representation learning across medical images and reports, *arXiv preprint arXiv:1811.08615* (2018).
24. Hassani K and Khasahmadi AH, Contrastive multi-view representation learning on graphs, *arXiv preprint arXiv:2006.05582* (2020).
25. Hotelling H, Relations between two sets of variates (Springer, 1992) pp. 162–190.
26. Li Y, Yang M and Zhang ZM, A survey of multi-view representation learning, *IEEE Transactions on Knowledge and Data Engineering* (2018).
27. Deng C, Chen Z, Liu X, Gao X and Tao D, Triplet-based deep hashing network for cross-modal retrieval, *IEEE Transactions on Image Processing* 27, 3893 (2018). [PubMed: 29993656]
28. Koch G, Zemel R and Salakhutdinov R, Siamese neural networks for one-shot image recognition, *ICML deep learning workshop 2* (2015).
29. Filzen TM, Kutchukian PS, Hermes JD, Li J and Tudor M, Representing high throughput expression profiles via perturbation barcodes reveals compound targets, *PLoS computational biology* 13, p. e1005335 (2017). [PubMed: 28182661]
30. Hoffer E and Ailon N, Deep metric learning using triplet network, *International Workshop on Similarity-Based Pattern Recognition*, 84 (2015).
31. Wu C-Y, Manmatha R, Smola AJ and Krahenbuhl P, Sampling matters in deep embedding learning, *Proceedings of the IEEE International Conference on Computer Vision*, 2840 (2017).
32. Weininger D, Weininger A and Weininger JL, Smiles. 2. algorithm for generation of unique smiles notation, *Journal of chemical information and computer sciences* 29, 97 (1989).
33. Landrum Get al., Rdkit: Open-source cheminformatics (2006).
34. Rogers D and Hahn M, Extended-connectivity fingerprints, *Journal of chemical information and modeling* 50, 742 (2010). [PubMed: 20426451]
35. Hodos R, Zhang P, Lee H-C, Duan Q, Wang Z, Clark NR, Maayan A, Wang F, Kidd B, Hu Jet al., Cell-specific prediction and application of drug-induced gene expression profiles, *Pacific Symposium on Biocomputing23* (2017).
36. McDermott M, Wang J, Zhao WN, Sheridan SD, Szolovits P, Kohane I, Haggarty SJ and Perlis RH, Deep learning benchmarks on 11000 gene expression data, *IEEE/ACM transactions on computational biology and bioinformatics* (2019).
37. Pedregosa F, Varoquaux G, Gramfort A, Michel V, Thirion B, Grisel O, Blondel M, Prettenhofer P, Weiss R, Dubourg Vet al., Scikit-learn: Machine learning in python, *Journal of machine learning research* 12, 2825 (2011).

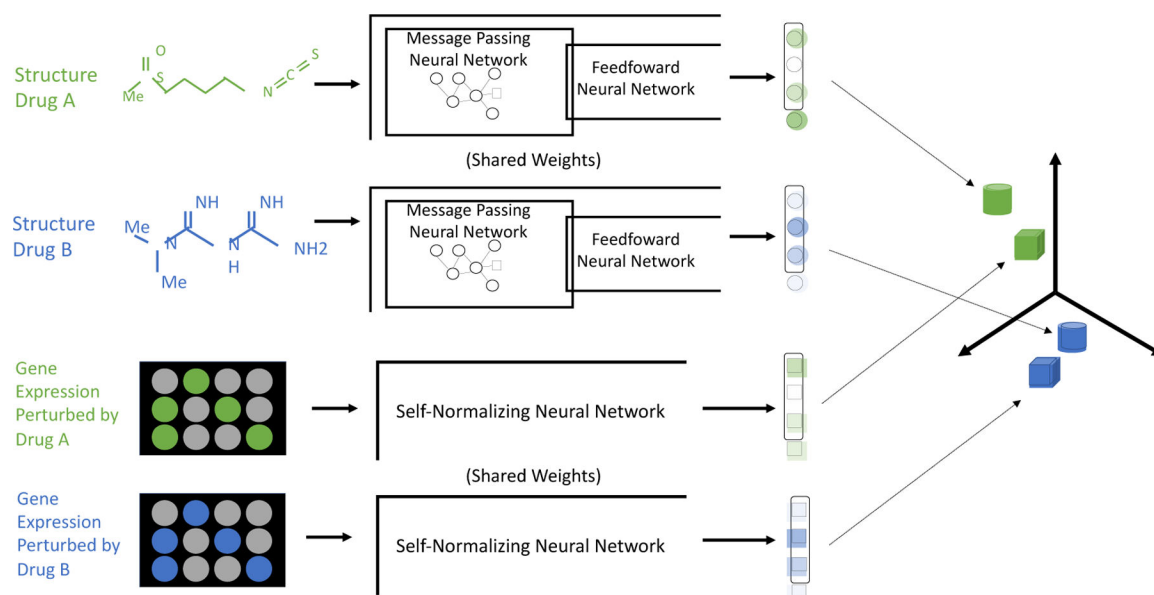
38. Bergstra J, Yamins D and Cox DD, Making a science of model search: Hyperparameter optimization in hundreds of dimensions for vision architectures, Proceedings of the 30th International Conference on International Conference on Machine Learning - Volume 28, p. II15I123 (2013).
39. Felix R, Kumar VB, Reid I and Carneiro G, Multi-modal cycle-consistent generalized zero-shot learning, Proceedings of the European Conference on Computer Vision (ECCV), 21 (2018).
40. McDermott MBA, Yan T, Naumann T, Hunt N, Suresh H, Szolovits P and Ghassemi M, Semi-Supervised Biomedical Translation with Cycle Wasserstein Regression GANs, p. 8 (2018).
41. Ying Z, Bourgeois D, You J, Zitnik M and Leskovec J, Gnnexplainer: Generating explanations for graph neural networks, Advances in neural information processing systems, 9244 (2019).

Author Manuscript


Author Manuscript

Author Manuscript

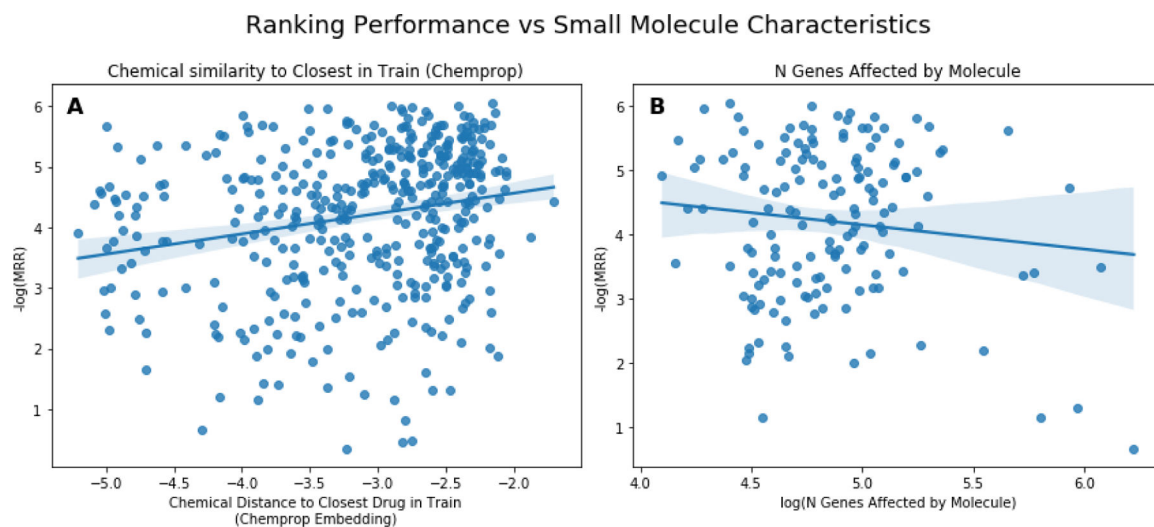
Author Manuscript



**Fig. 1.** Our representation learning method. Neural networks are trained to embed gene expression profiles close to the small molecule structures that induce them. Given a cross-modal alignment, gene expression signatures can be used as queries to rank chemical structures by their likelihood to induce such a signature.



**Fig. 2.** Precision Recall curves for drug identification given gene expression signatures, across various baselines (dashed lines), oracles (dotted lines) and our model (solid lines).

**Fig. 3.**

Left: Performance (lower is better) vs. structural distance to the nearest compound in the training set. This plot demonstrates that compounds more structurally dissimilar to the train set show mildly worse performance than those that are more similar. See Appendix Figure A2 for analogous plots for four additional measures of distance from the training set. Right: Average Performance vs. # of Genes differentially expression following treatment with the molecule of interest, showing that compounds that have broader transcriptomic effects are better retrieved by this method.

**Table 1.**

IR metrics across various configurations of the model/baselines. ‘Chemprop’ refers to pretrained model from Yang et al.<sup>6</sup> ‘Aligned’ indicates representations learned from our method (see Section 3.3). MR=median rank, MRR=mean reciprocal rank, H@K=Hit/Recall at K.

GE	Chemical	Method	MR	MRR	H@10	H@100
Raw	Morgan FP	D-NN	206	0.025	0.037	0.240
Raw	Chemprop	D-NN	211	0.025	0.035	0.254
Raw	Aligned	D-NN	189	0.033	0.047	0.290
Aligned	Morgan FP	D-NN	214	0.025	0.041	0.256
Aligned	Chemprop	D-NN	196	0.022	0.037	0.278
Aligned	Aligned	D-NN	137	0.039	0.072	0.402
Raw	Morgan FP	CCA	180	0.027	0.045	0.303
Raw	Chemprop	CCA	184	0.024	0.040	0.294
Raw	Aligned	CCA	134	0.039	0.076	0.412
Aligned	Morgan FP	CCA	177	0.027	0.050	0.319
Aligned	Chemprop	CCA	163	0.028	0.051	0.334
Aligned	Aligned	CCA	130	<b>0.048</b>	<b>0.093</b>	0.425
Aligned	Aligned	Corr	<b>126</b>	0.042	0.085	<b>0.432</b>

**Table 2.**

IR metrics for the various oracle methods.

<b>Oracle Model</b>	<b>MR</b>	<b>MRR</b>	<b>H@10</b>	<b>H@100</b>
Oracle - Train Cell Lines (Raw GE)	138	0.057	0.101	0.400
Oracle - Train Cell Lines (Embed)	110	0.064	0.128	0.466
Oracle - Test Cell Line (Raw GE)	82	0.076	0.147	0.565
Oracle - Test Cell Line (Embed)	79	0.093	0.160	0.568
Our Approach	126	0.042	0.085	0.432

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript