
Learning Deep Attribution Priors Based On Prior Knowledge

Ethan Weinberger

Paul G. Allen School of Computer Science
University of Washington
Seattle, WA 98195
ewein@cs.washington.edu

Joseph D. Janizek

Paul G. Allen School of Computer Science
University of Washington
Seattle, WA 98195
jjanizek@cs.washington.edu

Su-In Lee

Paul G. Allen School of Computer Science
University of Washington
Seattle, WA 98195
suinlee@cs.washington.edu

Abstract

Feature attribution methods, which explain an individual prediction made by a model as a sum of *attributions* for each input feature, are an essential tool for understanding the behavior of complex deep learning models. However, ensuring that models produce meaningful explanations, rather than ones that rely on noise, is not straightforward. Exacerbating this problem is the fact that attribution methods do not provide insight as to *why* features are assigned their attribution values, leading to explanations that are difficult to interpret. In real-world problems we often have sets of additional information for each feature that are predictive of that feature’s importance to the task at hand. Here, we propose the *deep attribution prior* (DAPr) framework to exploit such information to overcome the limitations of attribution methods. Our framework jointly learns a relationship between prior information and feature importance, as well as biases models to have explanations that rely on features predicted to be important. We find that our framework both results in networks that generalize better to out of sample data and admits new methods for interpreting model behavior.

1 Introduction

Recent advances in machine learning have come in the form of complex models that humans struggle to interpret. In response to the black-box nature of these models, a variety of recent work has focused on model interpretability [11]. One particular line of work that has gained much attention is that of feature attribution methods [22, 37, 28, 3, 33]. Given a model and a specific prediction made by that model, these methods assign a numeric value to each input feature, indicating how important that feature was for the given prediction (Figure 1a). A variety of such methods have been proposed, and previous work has focused on how such methods can be used to gain insight into model behavior in applications where model trust is critical [43, 31]

Given a set of attributions, a natural question is *why* a feature was assigned a specific attribution value. In some settings it is easy for a human to evaluate the sensibility of attributions; for example, in image classification problems we can overlay attribution values on the original image. However, in many other domains we do not have the ability to assess the validity of attribution values so

easily. Recent work [30, 7, 29] has attempted to address this problem by regularizing model training so that both model predictions and explanations agree with prior human knowledge. While such regularization methods do lead to noticeably different attributions, they suffer from two shortcomings. First, they require a human expert with prior knowledge about a given problem domain to construct a domain-specific regularization function. Second, after a set of feature attribution values is produced for a given prediction, we still do not have a human-interpretable way to understand why that set of values was produced beyond trusting in the regularization procedure.

At the same time, the overparameterization of deep learning models makes them prone to overfitting to noise. This limitation has stalled the adoption of deep models in domains where data acquisition is difficult, such as many areas of medicine [5]. As such, methods that encourage deep models to learn more generalizable representations, even when sample sizes are small, are in high demand.

In many real-world tasks we have access to sets of information about each feature that characterize the feature. We refer to such sets of information as *meta-features*. Meta-features can potentially encode information on a feature’s relative importance across all samples to the task at hand. For example, we can consider a task where we seek to predict a user’s rating for a new movie based on their previous ratings for other movies. In this case we would expect ratings for movies with similar meta-feature values to be more relevant than those for dissimilar films. Potential meta-features to capture the similarity between features in this task could include movie genre and director, among others. By using meta-features to bias models to focus on more relevant features when making predictions, we may be able to achieve both greater model accuracy and interpretability. For a given problem, while a practitioner may have an idea of potential meta-features that are relevant to the problem, it is far less likely that they know *a priori* a specific relationship between these meta-features and global feature importance. As a first step towards mitigating this problem, we propose the *deep attribution prior* (DAPr) framework for training deep neural networks that simultaneously learns a relationship between meta-features and global feature importance values, and biases the prediction model to have local explanations that broadly agree with this learned relationship (Figure 1b). For clarity, in the remainder of the text we use feature importance specifically to refer to a feature’s global relevance to a task across samples, and attribution value to refer to its local relevance for a specific prediction.

We apply our method to a synthetic dataset and two real-world biological datasets. Empirically we find that models trained using the DAPr framework achieve better performance on tasks where training data is limited. Furthermore, we demonstrate that such learned meta-feature to feature importance relationships can be leveraged to obtain more insight into the behavior of deep models.

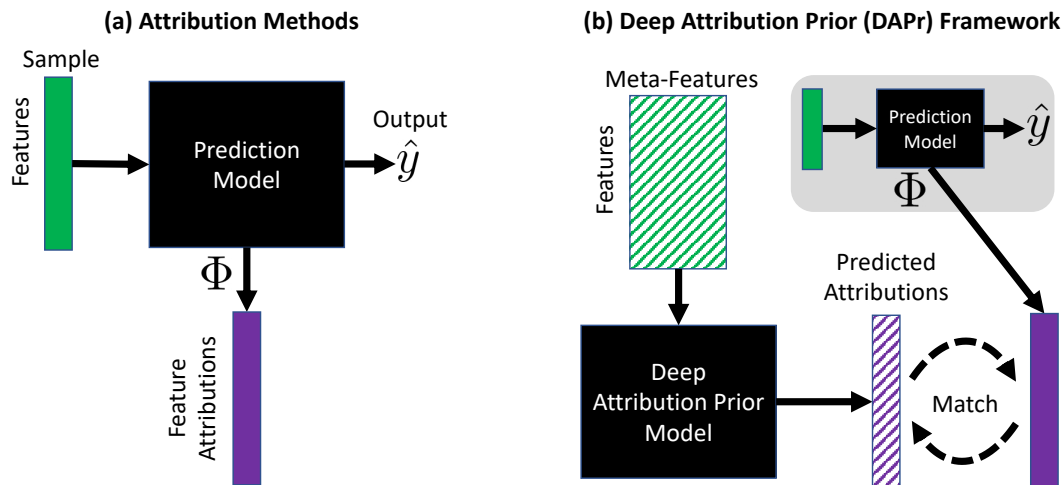


Figure 1: (a) An attribution method Φ is used to explain the decision of an arbitrary black-box model. (b) Overview of the DAPr framework. In addition to training a prediction model, we train a second prior model to use meta-features to predict global feature importance values. We also add a penalty term to our prediction model’s loss function that encourages the values of local attributions to be similar to predicted global importance values.

2 Related Work

Prior knowledge has long been used to guide the design of new machine learning algorithms. In a deep learning context, domain-specific knowledge has been essential to the design of model architectures that have achieved breakthroughs in performance. For example, LeCun et al. [17] took advantage of the smoothness of images to design a specialized neural network architecture for vision problems. Similarly, Bahdanau et al. [4] introduced the attention mechanism to alleviate problems with memorizing long source sentences in neural machine translation.

A line of recent work has studied using prior knowledge to bias the training of previously-existing network architectures so that their *explanations* align with prior knowledge. Such work [7, 30, 29] has focused on training networks by penalizing differences between local feature importance, as measured by an attribution method, and a function constructed by a human domain expert.

One particular form of prior knowledge is that of meta-features. Previous work has studied using meta-features to constrain learning algorithms to learn more generalizable models. Lee et al. [18] devised an algorithm for linear models that uses meta-features to transfer knowledge across related tasks, assuming that a feature’s weight is a function of its meta-feature values. Krupka and Tishby [15] proposed a framework for training support vector machines using meta-features to define a prior on the prediction model’s weights. More recently Lee et al. [19] proposed the MERGE algorithm, which trains a linear model so that its weights match the output of a second linear model that learns a relationship between feature weights and meta-feature values. The two models are iteratively optimized until convergence. However, to our knowledge DAPr is the first method to bias the training of neural networks using prior knowledge in the form of meta-features.

3 Deep Attribution Priors

In this section we introduce a formal definition of a deep attribution prior by extending that of the attribution prior proposed in Erion et al. [7]. Let $X \in \mathbb{R}^{n \times p}$ denote a training dataset consisting of n samples, each of which has p features. Similarly, let $Y \in \mathbb{R}^n$ denote labels for the samples in our training dataset. In a standard model training procedure, we wish to find the model \hat{f} in some model class \mathcal{F} that minimizes the average prediction loss on our dataset, as determined by some loss function \mathcal{L} . In order to avoid overfitting on training data, a regularization function Ω' on the prediction model’s parameters θ_f is often included, giving the equation

$$\hat{f} = \operatorname{argmin}_{f \in \mathcal{F}} \frac{1}{n} \sum_{x,y} \mathcal{L}(f(x), y) + \lambda' \Omega'(\theta_f),$$

where λ' is a scalar hyperparameter controlling the strength of the regularization. For a given feature attribution method $\Phi(\theta, x)$, an attribution prior is defined as some function $\Omega: \mathbb{R}^p \rightarrow \mathbb{R}$ that assigns scalar-valued penalties to the feature attributions for f with input x . This leads us to the following objective:

$$\hat{f} = \operatorname{argmin}_{f \in \mathcal{F}} \frac{1}{n} \sum_{x,y} \mathcal{L}(f(x), y) + \lambda \sum_x \Omega(\Phi(\theta_f, x)).$$

Now suppose that each feature in X is associated with some scalar-valued meta-feature. We can represent these meta-feature values as a vector $m \in \mathbb{R}^p$, where the i -th entry in m represents the value of our meta-feature corresponding to the i -th feature in the prediction problem. If, in some hypothetical scenario, we knew beforehand that the values in m corresponded to each feature’s global importance to the task at hand, a natural choice for Ω would be a penalty on the difference between our model’s feature attributions and the values in m , giving us

$$\hat{f} = \operatorname{argmin}_{f \in \mathcal{F}} \frac{1}{n} \sum_{x,y} \mathcal{L}(f(x), y) + \lambda \sum_x \|\Phi(\theta_f, x) - m\|_1, \quad (1)$$

where $\|\cdot\|_1$ is the L_1 norm. We note that this function penalizes a difference between local attributions (Φ), and global feature importance. While such a prior model may not make sense for some domains

(e.g. image classification tasks), previous work has empirically demonstrated benefits in prediction performance from regularizing attributions with a global prior in contexts where we might expect local attributions to be more correlated with a global prior (e.g. tabular biomedical datasets). Indeed, the results of Erion et al. [7] on gene expression data demonstrated improved generalization when regularizing deep networks such that their local attributions were close to pre-computed global importance values. Having such an m that is known beforehand to be perfectly predictive of feature importance is unrealistic in most settings. However, suppose instead that we have some meta-feature matrix $M \in \mathbb{R}^{p \times k}$, corresponding to each of the p features in the prediction problem having k meta-features. If we were to assume that there is a linear relationship between a feature’s meta-feature values and that feature’s importance, we can modify Equation 1 to achieve

$$\hat{f} = \operatorname{argmin}_{f \in \mathcal{F}} \frac{1}{n} \sum_{x,y} \mathcal{L}(f(x), y) + \lambda \sum_x \|\Phi(\theta_f, x) - \underbrace{\beta_1 M_{*,1} + \dots + \beta_k M_{*,k}}_{\text{Linear attribution prior}}\|_1,$$

where $M_{*,i}$ refers to the i -th column of M , and $\beta \in \mathbb{R}^k$ is a vector of hyperparameters that we can optimize over. Finding an optimal set of hyperparameters gives us a model that captures the meta-feature to feature importance relationship for a given problem. Assuming a linear relationship is restrictive, and likely fails to capture many of these relationships. Instead, we can take advantage of the universal function approximation capabilities of neural networks and replace our linear model with a deep one $\hat{g}: \mathbb{R}^k \rightarrow \mathbb{R}$ from some class of networks \mathcal{G} . We define such a model to be a deep attribution prior (DAPr). As either the number of meta-features k increases or \mathcal{G} becomes more complex, thereby leading to more model parameters, learning the parameters of a DAPr by treating them as hyperparameters becomes computationally prohibitive. Instead, we jointly learn a prediction model and DAPr pair (\hat{f}, \hat{g}) by optimizing the following objective:

$$(\hat{f}, \hat{g}) = \operatorname{argmin}_{f \in \mathcal{F}, g \in \mathcal{G}} \frac{1}{n} \sum_{x,y} \mathcal{L}(f(x), y) + \lambda \sum_x \|\Phi(\theta_f, x) - G(M)\|_1,$$

where we define $G(M)$ to be the vector in \mathbb{R}^p resulting from feeding each row of M into g (i.e., $G(M)$ is the vector of predicted importance values for all of our features based on each feature’s meta-feature values). We are able to approximate a solution to this problem by alternating between one step of optimizing \hat{f} , and one step of optimizing \hat{g} via a gradient descent-based optimization method. We find empirically that, so long as \hat{f} and \hat{g} change slowly enough, this procedure accomplishes the following two goals:

- Learn a relationship between provided meta-features and feature importance for a given prediction task.
- Bias prediction models to rely more heavily on the features deemed important by this learned meta-feature to feature importance relationship.

We show in Section 4 that introducing such bias in the training procedure improves the performance of deep neural network models in multiple settings where data is limited. Furthermore, we demonstrate in Section 5 that we can exploit these learned meta-feature to feature importance relationships to achieve new insights into our prediction models’ behavior. For the experiments in the remainder of the text we use Expected Gradients (EG) [7] for our attribution method Φ because it has a low computational cost and resulted in superior performance as compared to other gradient-based attribution methods; we refer the reader to the Supplement for details on EG and additional experiments examining performance with other attribution methods.

4 Training with Deep Priors Improves Accuracy in Low-Data Settings

4.1 Two Moons Classification With Nuisance Features

We first evaluate the performance of multi-layer perceptrons (MLPs) with two hidden layers trained using the DAPr framework on the two moons with nuisance features task as introduced in Yamada et al. [41]. In particular we seek to understand whether our framework biases MLPs towards relying

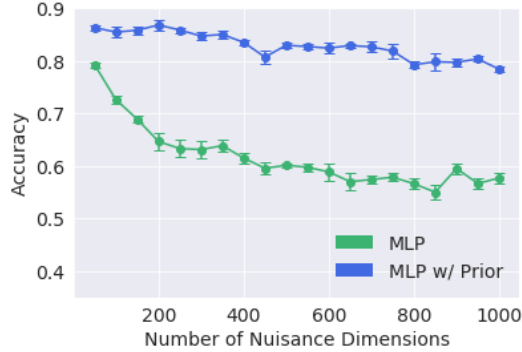


Figure 2: Classification accuracy (mean and standard error) for the two moons with nuisance features task of MLPs trained with and without the DAPr framework.

on informative features. In this experiment we construct a dataset based on two two-dimensional moon shape classes, concatenated with noisy features. For a given data point the first two coordinates x_1, x_2 are drawn by adding noise drawn from an $\mathcal{N}(0, 0.1)$ distribution to a point from one of two nested half circles. The remaining coordinates x_3, \dots, x_p are drawn from an $\mathcal{N}(0, 1)$ distribution. Our goal is to classify points as originating from one half-circle or the other.

For each of our experiments, we generate five datasets consisting of 1000 samples. For each dataset we use 20% of the dataset for model training, and divide the remaining points evenly into testing and validation sets, giving a final 20%-40%-40% train-test-validation split. We vary the number of nuisance features from 50 to 1000 in increments of 50. For a given number of features p we construct a meta-feature matrix $M_{\text{moons}} \in \mathbb{R}^{p \times 2}$, where the i -th row contains the i -th feature’s mean and standard deviation. For our prior we train a linear model that attempts to predict feature importance based on that feature’s mean and standard deviation. The number of units in the hidden layers of our MLPs depends on the number of features p ; for a given p the first hidden layer has $\lfloor p/2 \rfloor$ units, and the second has $\lfloor p/4 \rfloor$ units.

We optimize our models using Adam [14] with early stopping, and all hyperparameters are chosen based on performance on our validation sets (see Supplement for additional details). We report our results in Figure 2. We find that our MLPs trained with prior models using M_{moons} are far more robust to the addition of noisy features than MLPs that do not take advantage of meta-features. This result continues to hold even as the number of features in the dataset grows far beyond the number of training points. This phenomenon indicates that training an MLP with an appropriate prior model is able to bias the MLP towards learning meaningful relationships rather than overfitting to noise.

4.2 Alzheimer’s Disease Biomarker Prediction

Alzheimer’s disease (AD) is a slowly progressing neurodegenerative disorder affecting an estimated 5.7 million Americans [2]. Due to a rapidly aging population, this number is expected to increase to 13.8 million by 2050, and thus AD poses an increasing threat to healthcare systems. While AD has well-known biomarkers in the form of amyloid- β protein deposits and misfolded tau proteins, the genetic drivers for these markers remain unclear [27].

In this experiment we explore how MLPs trained with the DAPr framework perform at using bulk RNA-seq data to predict amyloid- β levels measured using immunohistochemistry. The data for this experiment comes from multiple datasets available on the Accelerating Medicines Partnership Alzheimer’s Disease Project (AMP-AD) portal¹; in particular we make use of the Adult Changes in Thought (ACT) [24], Mount Sinai Brain Bank (MSBB), and Religious Orders Study/Memory and Aging Project (ROSMAP) [1] datasets. All together these studies provide us with $n = 1742$ labelled gene expression samples. For prior information we gather a set of AD driver meta-features for each gene consisting of a binary variable indicating the presence of copy number variations, a second binary variable indicating whether the gene is a known regulator of other genes, methylation

¹<https://adknowledgeportal.synapse.org>

Table 1: Performance (MSE \pm SE) of MLPs trained with DAPr using disease-specific driver features vs. baseline models. * indicates the use of meta-features during training.

Model	AD Biomarker Prediction	AML Drug Response
LASSO	1.19 \pm 0.35	1.031 \pm 0.051
MLP	0.78 \pm 0.15	0.938 \pm 0.106
MLP (L1 Reg.)	0.74 \pm 0.12	0.838 \pm 0.070
MLP (L2 Reg.)	0.75 \pm 0.14	0.871 \pm 0.075
MLP w/ Drivers*	1.059 \pm 0.24	1.158 \pm 0.117
MLP w/ DAPr (noise)*	0.82 \pm 0.15	0.915 \pm 0.079
MERGE [19]*	1.07 \pm 0.21	0.902 \pm 0.082
MLP w/ DAPr (drivers)*	0.67 \pm 0.13	0.786 \pm 0.065

values, and node strength in a gene-gene interaction graph. We refer the reader to the Supplement for additional details on data collection and preprocessing. Using this prior information we construct a meta-feature matrix $M_{AD} \in \mathbb{R}^{p \times 4}$, where p is the number of genes for which we have both expression data and driver feature values. In this experiment $p = 12, 326$.

We compare our framework to multiple baselines and report our results in Table 1. To understand how incorporating meta-features into the training process improves model performance, we compare against LASSO [38], and MLPs trained without prior information. We also train MLPs with L_1 and L_2 regularization on their weights to illustrate how much of our performance gains comes from using prior knowledge specifically, as opposed to regularization in general. To demonstrate the benefits of using DAPr in particular to integrate meta-features into the training process, we compare against MLPs that naively make use of meta-features by simply treating them as additional input features. In addition, to confirm the informativeness of our specific meta-features, we compare training MLPs with DAPr models using AD driver features to MLPs also trained with DAPr models, but for which the prior information is composed solely of Gaussian noise. Finally, we compare our framework to the previous state of the art method for biasing model training using meta-features, MERGE [19], which biases a linear model’s weights to match the output of a second linear model trained to predict weights from meta-feature values. All models are evaluated on five splits of the data into training, validation, and test sets. For each split we use 60% of the data for model training, 20% for validation, and 20% for testing. All prediction model MLPs have two hidden layers with 512 and 256 units respectively, and for our DAPr models we use MLPs with one hidden layer containing four units. For fairness we train all models using only the p genes for which we have both expression data and meta-feature values. We refer the reader to the Supplement for further details on model training and evaluation.

We find that introducing meta-features into the training process can lead to better performance for both linear models and deep ones. However, care is required when using meta-features to train deep models; our naive method for training MLPs with meta-features leads to a degradation in performance. Our results also make clear the limitations of the linear models in MERGE, as even standard MLPs outperform it. Unsurprisingly, MLPs trained with DAPr models using noise meta-features perform slightly worse than standard MLPs. On the other hand, we find that MLPs trained with DAPr models using AD driver features see a 13.5% decrease in MSE as compared to standard MLPs. While we find minor boosts in performance by training MLPs with L_1 and L_2 regularization, these improvements are far less than that of the MLPs with informative DAPr models. These results further indicate that DAPr models can bias MLPs towards learning meaningful nonlinear relationships not captured by simpler models even in high-dimensional, low sample size settings.

4.3 AML Drug Response Prediction

Acute myeloid leukemia (AML) is a form of blood cancer characterized by the rapid buildup of abnormal cells in the blood and bone marrow that interfere with the function of healthy blood cells. The disease has a very poor prognosis, leading to death in approximately 80% of patients, and it is the leading cause of leukemia-related deaths in the United States [40]. This phenomenon is likely due in large part to the heterogeneity of AML; different AML patients have varying responses to the same treatment regimens, and methods for better matching patients to drugs are in high demand.

In this experiment we explore how training an MLP with a DAPr model affects performance on an AML drug response prediction task. Our data comes from the Beat AML dataset, containing gene expression data and drug sensitivity measures for tumors from 572 patients [39]. In our analysis we focus on the drug Dasatinib, which had data from the most patients ($n = 217$). For meta-features we use the publicly available AML driver features from Lee et al. [19]: mutation frequencies, a binary variable indicating the presence of copy number variations, a second binary variable for whether the gene is a known regulator of other genes, gene expression hubness, and methylation values.

For this experiment we compare MLPs trained with DAPr models to the same baselines as in Section 4.2. We evaluate each model on five splits of the data. For each split, 80% of the data is used for training, while the remaining 20% is divided evenly into validation and test sets. We use the same prediction model MLP architecture as in Section 4.2. For our DAPr models we use MLPs with two hidden layers, of five and three units, to capture the relationship between meta-features and feature importance. We refer the reader to the Supplement for details on hyperparameter tuning.

Once again our results demonstrate that incorporating informative meta-features into the model training process can result in boosts in performance for both linear models and deep ones. While in this case MERGE performs on par with standard MLPs, our best performing models once again are MLPs trained with DAPr models.

5 Deep Attribution Priors Admit New Insights into Deep Models

While feature attribution methods provide users with a sense of the “relevant” features for a given prediction, they lack a way to contextualize the attribution values in a human-interpretable way. This shortcoming can lead to a false sense of security when employing such methods, even though previous work has demonstrated their potential to produce explanations known to be nonsensical based on prior human knowledge [42, 13, 10]. However, we can take a step towards alleviating this problem by probing a DAPr model \hat{g} to understand how meta-features drive changes in its predicted global feature importance values. In the case where \hat{g} is a linear model, we can simply use the model’s weight coefficients to get an explanation for a given predicted importance value. When \hat{g} is a deep model, uncovering such explanations is not as straightforward. However, in this section, using the AML drug response prediction problem from Section 4.3 as a case study, we demonstrate multiple methods for probing DAPr models to obtain insights into their predictions. Using such methods we find that our DAPr model from Section 4.3 independently learns meta-feature to gene importance relationships that agree with prior biological knowledge.

5.1 Attribution Explanations For Understanding Meta-Feature to Gene Relationships

Given a potentially complex DAPr model \hat{g} and vector of meta-feature values m_i for the i -th feature in a prediction problem, we would like to understand how the values of m_i relate to that feature’s predicted global importance value $\hat{g}(m_i)$. To do so we can apply an attribution method Φ to \hat{g} , thereby generating a second order explanation $\Phi(\theta_{\hat{g}}, m_i)$, explaining the i -th feature’s predicted global importance in terms of human-interpretable meta-features. We apply EG as our Φ to the DAPr model from Section 4.3, and visualize our results in Figure 3.

Remarkably, we find that our DAPr model rediscovers relationships in line with prior biological knowledge. We can see in Figure 3 that high predicted gene importance values are explained mostly by expression hubness. This trend is consistent with prior knowledge, as expression hubness has been suspected to drive events in cancer [21]. Furthermore, we observe that, for a small number of genes, mutation appears as a noticeable factor in their importance explanations. This phenomenon also agrees with prior knowledge, as mutations in the genes CEPBA, FLT3, and ELF4 are suspected to play a role in the variability of drug response between AML patients [20, 8, 34, 6, 36]. We also note that our framework, on its own, selected for relevant meta-features among those provided. This behavior validates that DAPr can reduce the domain knowledge requirement of the original attribution prior framework; rather than needing a specific meta-feature to feature important relationship up front, DAPr can select for relevant meta-features out of those provided.

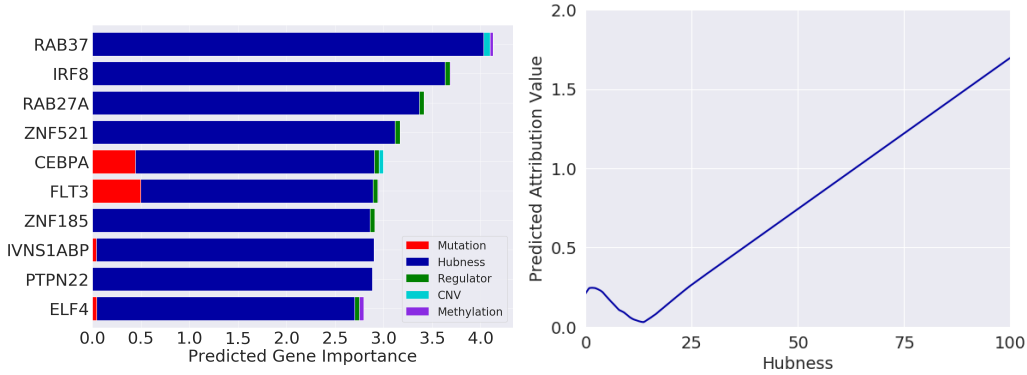


Figure 3: (Left) Explanations of feature importance for most important genes as ranked by absolute predicted importance. Bar color proportional to absolute value of that meta-feature’s contribution to the predicted importance value as determined by EG. (Right) Deep attribution prior PDP for expression hubness.

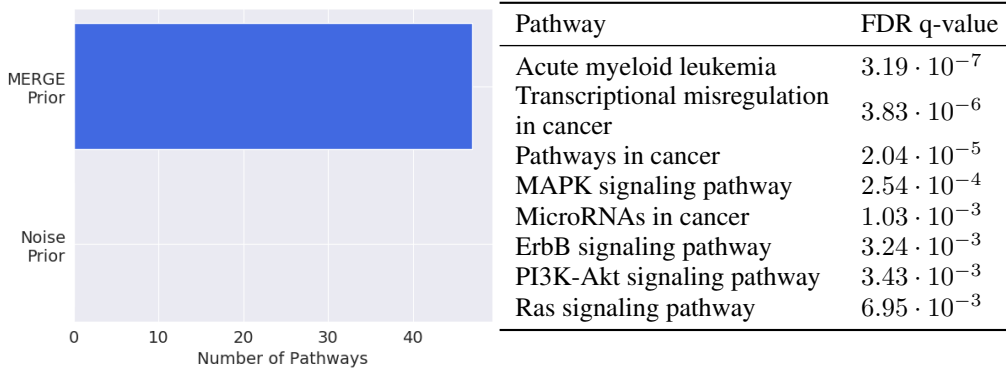


Figure 4: (Left) Number of pathways captured by DAPr trained using AML driver features vs. DAPr trained with noise. (Right) Sample of captured pathways believed to be relevant to AML.

5.2 Partial Dependence Plots Capture Nonlinear Meta-Feature to Feature Importance Relationships

We can further explore the nature of these learned meta-feature to feature importance relationships by constructing partial dependence plots [9] to visualize the marginal effect of particular meta-features on predicted global importance values. We display the results of doing so for expression hubness in Figure 3. We find that our DAPr model captures a non-linear relationship between hubness and predicted global importance. This relationship broadly agrees with prior knowledge; after an initial buildup increases in hubness are strongly associated with a gene’s relevance to drug response. Our other meta-features also exhibited nonlinear relationships with predicted feature importance, and we refer the reader to the Supplement for additional figures. The nonlinearities captured by our DAPr model indicate that more complex models may better capture the relationship between a gene’s meta-features and its potential to drive events in AML than e.g. the linear models used by MERGE.

5.3 Deep Attribution Priors Capture Relevant Gene Pathways

To further confirm that the relationships captured by our DAPr model match biological intuition, we perform Gene Set Enrichment Analysis [35] to see if the top genes as ranked by our prior were enriched for membership in any biological pathways. For comparison we use the number of pathways captured by a DAPr model trained on noise as a baseline. In our analysis we use the top 200 genes as ranked by both prior models, and we use the Enrichr [16] library to check for membership in the Kyoto Encyclopedia of Genes and Genomes [12] 2019 pathways list. While the top genes as ranked by our noise DAPr model are not significantly enriched for membership in *any* pathways

after FDR correction, our AML driver DAPr model captures many pathways, of which multiple are already believed to be associated with AML [23, 32, 25, 26]. This result further indicates that our deep attribution prior is capturing meaningful meta-feature to gene importance relationships.

6 Discussion

In this work we introduce the DAPr framework for biasing the training of neural networks by incorporating prior information about the features used for a prediction task. Unlike other feature attribution based penalty methods, our framework merely requires the presence of prior information in the form of meta-features, rather than specific rules hand-crafted by a domain expert. In our experiments we find that jointly learning prediction models and DAPr models leads to increased performance on prediction tasks in settings with limited data. Furthermore, we demonstrate that the DAPr framework admits new methods for understanding deep models. Using such methods we demonstrate that our prior models independently learn meta-feature to feature importance relationships that agree with prior knowledge. The DAPr framework provides a broadly applicable method for incorporating prior knowledge in the form of meta-features into the training of deep neural networks, and we believe that it is a valuable tool for learning in low sample size, trust-critical domains.

Broader Impact

DAPr can be applied to a wide variety of problems for which prior knowledge is available about a dataset’s individual features. In our work we focus on applying our method to a synthetic dataset and two real-world medical datasets, though it should be easily extendable to other problem domains.

As discussed in the introduction, a major barrier to the adoption of modern machine learning techniques in real-world settings is that of *trust*. In high-stakes domains, such as medicine, practitioners are wary of replacing human judgement with that of black box algorithms, even if the black box consistently outperforms the human in controlled experiments. This concern is well-founded, as many high-performing systems developed in research environments have been found to overfit to quirks in a particular dataset, rather than learn more generalizable patterns. In our work we demonstrate that the DAPr framework does help deep networks generalize to our test sets when sample sizes are limited. While these results are encouraging, debugging model behavior in the real world, where data cannot simply be divided into training and test sets, is a more challenging problem. Feature attribution methods are one potential avenue for debugging models; however, while it may be easy to tell from a set of attributions if e.g. an image model is overfitting to noise, it would be much more difficult for a human to determine that a model trained on gene expression data was learning erroneous patterns simply by looking at attributions for individual genes. By learning to explain a given feature’s global importance using meta-features, we believe that DAPr can provide meaningful insights into model behavior that can help practitioners debug their models and potentially deploy them in real-world settings.

Nevertheless, we recognize the potential downsides with the adoption of complex machine learning interpretability tools. Previous results have demonstrated that interpretability systems can in fact lead users to have *too much* trust in models when a healthy dose of skepticism would be more appropriate. More research is needed to understand how higher-order explanation tools like DAPr influence user behavior to determine directions for future work.

Acknowledgments and Disclosure of Funding

This work was supported by funding from the National Science Foundation [CAREER DBI-1552309 and DBI-1759487], American Cancer Society [127332-RSG-15-097-01-TBG], and National Institutes of Health [R35 GM 128638 and R01 NIA AG 061132]. The results published in this work are in part based on data obtained from the Alzheimer’s Disease Knowledge Portal (<https://adknowledgeportal.synapse.org/>) operated by Sage Bionetworks. The results published here are partially based upon data generated by the Cancer Target Discovery and Development (CTD2) Network (<https://ocg.cancer.gov/programs/ctd2/data-portal>) established by the National Cancer Institute’s Office of Cancer Genomics. We thank the members of the Lee Lab and the anonymous NeurIPS reviewers for their feedback that significantly improved this work.

References

- [1] David A Bennett, Julie A Schneider, Zoe Arvanitakis, and Robert S Wilson. Overview and findings from the religious orders study. *Current Alzheimer Research*, 9(6):628–645, 2012.
- [2] Alzheimer’s Association. 2018 alzheimer’s disease facts and figures. *Alzheimer’s & Dementia*, 14(3):367–429, 2018.
- [3] Sebastian Bach, Alexander Binder, Grégoire Montavon, Frederick Klauschen, Klaus-Robert Müller, and Wojciech Samek. On pixel-wise explanations for non-linear classifier decisions by layer-wise relevance propagation. *PLoS one*, 10(7), 2015.
- [4] Dzmitry Bahdanau, Kyunghyun Cho, and Yoshua Bengio. Neural machine translation by jointly learning to align and translate. *arXiv preprint arXiv:1409.0473*, 2014.
- [5] David Chen, Sijia Liu, Paul Kingsbury, Sunghwan Sohn, Curtis B Storlie, Elizabeth B Habermann, James M Naessens, David W Larson, and Hongfang Liu. Deep learning and alternative learning strategies for retrospective real-world clinical data. *NPJ digital medicine*, 2(1):1–5, 2019.
- [6] Naval Daver, Richard F Schlenk, Nigel H Russell, and Mark J Levis. Targeting flt3 mutations in aml: review of current knowledge and evidence. *Leukemia*, 33(2):299–312, 2019.
- [7] Gabriel Erion, Joseph D Janizek, Pascal Sturmfels, Scott Lundberg, and Su-In Lee. Learning explainable models using attribution priors. *arXiv preprint arXiv:1906.10670*, 2019.
- [8] A Fasan, C Haferlach, T Alpermann, S Jeromin, V Grossmann, C Eder, S Weissmann, F Dicker, A Kohlmann, S Schindela, et al. The role of different genetic subtypes of cebpa mutated aml. *Leukemia*, 28(4):794, 2014.
- [9] Jerome H Friedman. Greedy function approximation: a gradient boosting machine. *Annals of statistics*, pages 1189–1232, 2001.
- [10] Yash Goyal, Uri Shalit, and Been Kim. Explaining classifiers with causal concept effect (cace). *arXiv preprint arXiv:1907.07165*, 2019.
- [11] Riccardo Guidotti, Anna Monreale, Salvatore Ruggieri, Franco Turini, Fosca Giannotti, and Dino Pedreschi. A survey of methods for explaining black box models. *ACM Comput. Surv.*, 51(5):93:1–93:42, August 2018. ISSN 0360-0300. doi: 10.1145/3236009. URL <http://doi.acm.org/10.1145/3236009>.
- [12] Minoru Kanehisa and Susumu Goto. Kegg: kyoto encyclopedia of genes and genomes. *Nucleic acids research*, 28(1):27–30, 2000.
- [13] Pieter-Jan Kindermans, Sara Hooker, Julius Adebayo, Maximilian Alber, Kristof T Schütt, Sven Dähne, Dumitru Erhan, and Been Kim. The (un) reliability of saliency methods. In *Explainable AI: Interpreting, Explaining and Visualizing Deep Learning*, pages 267–280. Springer, 2019.
- [14] Diederik P Kingma and Jimmy Ba. Adam: A method for stochastic optimization. *arXiv preprint arXiv:1412.6980*, 2014.
- [15] Eyal Krupka and Naftali Tishby. Incorporating prior knowledge on features into learning. In *Artificial Intelligence and Statistics*, pages 227–234, 2007.
- [16] Maxim V Kuleshov, Matthew R Jones, Andrew D Rouillard, Nicolas F Fernandez, Qiaonan Duan, Zichen Wang, Simon Koplev, Sherry L Jenkins, Kathleen M Jagodnik, Alexander Lachmann, et al. Enrichr: a comprehensive gene set enrichment analysis web server 2016 update. *Nucleic acids research*, 44(W1):W90–W97, 2016.
- [17] Yann LeCun, Léon Bottou, Yoshua Bengio, Patrick Haffner, et al. Gradient-based learning applied to document recognition. *Proceedings of the IEEE*, 86(11):2278–2324, 1998.
- [18] Su-In Lee, Vassil Chatalbashev, David Vickrey, and Daphne Koller. Learning a meta-level prior for feature relevance from multiple related tasks. In *Proceedings of the 24th international conference on Machine learning*, pages 489–496. ACM, 2007.

- [19] Su-In Lee, Safiye Celik, Benjamin A Logsdon, Scott M Lundberg, Timothy J Martins, Vivian G Oehler, Elihu H Estey, Chris P Miller, Sylvia Chien, Jin Dai, et al. A machine learning approach to integrate big data for precision medicine in acute myeloid leukemia. *Nature communications*, 9(1):42, 2018.
- [20] Liang-In Lin, Chien-Yuan Chen, Dong-Tsamn Lin, Woei Tsay, Jih-Luh Tang, You-Chia Yeh, Hwei-Ling Shen, Fang-Hsien Su, Ming Yao, Sheng-Yi Huang, et al. Characterization of cebpa mutations in acute myeloid leukemia: most patients with cebpa mutations have biallelic mutations and show a distinct immunophenotype of the leukemic cells. *Clinical cancer research*, 11(4):1372–1379, 2005.
- [21] Benjamin A Logsdon, Andrew J Gentles, Chris P Miller, C Anthony Blau, Pamela S Becker, and Su-In Lee. Sparse expression bases in cancer reveal tumor drivers. *Nucleic acids research*, 43(3):1332–1344, 2015.
- [22] Scott M Lundberg and Su-In Lee. A unified approach to interpreting model predictions. In *Advances in Neural Information Processing Systems*, pages 4765–4774, 2017.
- [23] Michele Milella, Steven M Kornblau, Zeev Estrov, Bing Z Carter, H el ene Lapillonne, David Harris, Marina Konopleva, Shourong Zhao, Elihu Estey, and Michael Andreeff. Therapeutic targeting of the mek/mapk signal transduction module in acute myeloid leukemia. *The Journal of clinical investigation*, 108(6):851–859, 2001.
- [24] Jeremy A Miller, Angela Guillozet-Bongaarts, Laura E Gibbons, Nadia Postupna, Anne Renz, Allison E Beller, Susan M Sunkin, Lydia Ng, Shannon E Rose, Kimberly A Smith, et al. Neuropathological and transcriptomic characteristics of the aged brain. *Elife*, 6:e31126, 2017.
- [25] Sophie Park, Nicolas Chapuis, J er ome Tamburini, Val erie Bardet, Pascale Cornillet-Lefebvre, Lise Willems, Alexa Green, Patrick Mayeux, Catherine Lacombe, and Didier Bouscary. Role of the pi3k/akt and mtor signaling pathways in acute myeloid leukemia. *haematologica*, 95(5): 819–828, 2010.
- [26] Emily J Pomeroy and Craig E Eckfeldt. Targeting ras signaling in aml: Ralb is a small gtpase with big potential. *Small GTPases*, pages 1–6, 2017.
- [27] Christiane Reitz. Alzheimer’s disease and the amyloid cascade hypothesis: a critical review. *International journal of Alzheimer’s disease*, 2012, 2012.
- [28] Marco Tulio Ribeiro, Sameer Singh, and Carlos Guestrin. Why should i trust you?: Explaining the predictions of any classifier. In *Proceedings of the 22nd ACM SIGKDD international conference on knowledge discovery and data mining*, pages 1135–1144. ACM, 2016.
- [29] Laura Rieger, Chandan Singh, W James Murdoch, and Bin Yu. Interpretations are useful: penalizing explanations to align neural networks with prior knowledge. *arXiv preprint arXiv:1909.13584*, 2019.
- [30] Andrew Slavin Ross, Michael C Hughes, and Finale Doshi-Velez. Right for the right reasons: Training differentiable models by constraining their explanations. *arXiv preprint arXiv:1703.03717*, 2017.
- [31] Rory Sayres, Ankur Taly, Ehsan Rahimy, Katy Blumer, David Coz, Naama Hammel, Jonathan Krause, Arunachalam Narayanaswamy, Zahra Rastegar, Derek Wu, et al. Using a deep learning algorithm and integrated gradients explanation to assist grading for diabetic retinopathy. *Ophthalmology*, 126(4):552–564, 2019.
- [32] Xuejing Shao, Yujia Liu, Yangling Li, Miao Xian, Qian Zhou, Bo Yang, Meidan Ying, and Qiaojun He. The her2 inhibitor tak165 sensitizes human acute myeloid leukemia cells to retinoic acid-induced myeloid differentiation by activating mek/erk mediated rara/stat1 axis. *Scientific reports*, 6:24589, 2016.
- [33] Avanti Shrikumar, Peyton Greenside, and Anshul Kundaje. Learning important features through propagating activation differences. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 3145–3153. JMLR. org, 2017.

- [34] Donald Small. Flt3 mutations: biology and treatment. *ASH Education Program Book*, 2006(1): 178–184, 2006.
- [35] Aravind Subramanian, Pablo Tamayo, Vamsi K Mootha, Sayan Mukherjee, Benjamin L Ebert, Michael A Gillette, Amanda Paulovich, Scott L Pomeroy, Todd R Golub, Eric S Lander, et al. Gene set enrichment analysis: a knowledge-based approach for interpreting genome-wide expression profiles. *Proceedings of the National Academy of Sciences*, 102(43):15545–15550, 2005.
- [36] Mary Ann Suico, Tsuyoshi Shuto, and Hirofumi Kai. Roles and regulations of the ets transcription factor elf4/mef. *Journal of molecular cell biology*, 9(3):168–177, 2017.
- [37] Mukund Sundararajan, Ankur Taly, and Qiqi Yan. Axiomatic attribution for deep networks. In *Proceedings of the 34th International Conference on Machine Learning-Volume 70*, pages 3319–3328. JMLR. org, 2017.
- [38] Robert Tibshirani. Regression shrinkage and selection via the lasso. *Journal of the Royal Statistical Society: Series B (Methodological)*, 58(1):267–288, 1996.
- [39] Jeffrey W Tyner, Cristina E Tognon, Daniel Bottomly, Beth Wilmot, Stephen E Kurtz, Samantha L Savage, Nicola Long, Anna Reister Schultz, Elie Traer, Melissa Abel, et al. Functional genomic landscape of acute myeloid leukaemia. *Nature*, 562(7728):526, 2018.
- [40] Justin Watts and Stephen Nimer. Recent advances in the understanding and treatment of acute myeloid leukemia. *F1000Research*, 7, 2018.
- [41] Yutaro Yamada, Ofir Lindenbaum, Sahand Negahban, and Yuval Kluger. Feature selection using stochastic gates. *arXiv preprint arXiv:1810.04247*, 2018.
- [42] Mengjiao Yang and Been Kim. Bim: Towards quantitative evaluation of interpretability methods with ground truth. *arXiv preprint arXiv:1907.09701*, 2019.
- [43] John R Zech, Marcus A Badgeley, Manway Liu, Anthony B Costa, Joseph J Titano, and Eric Karl Oermann. Variable generalization performance of a deep learning model to detect pneumonia in chest radiographs: A cross-sectional study. *PLoS medicine*, 15(11):e1002683, 2018.