MultiPLIER: a transfer learning framework reveals systemic features of rare autoimmune disease

Jaclyn N. Taroni[1,2,†], Peter C. Grayson[3], Qiwen Hu[1], Sean Eddy[4], Matthias Kretzler[4,5], Peter A. Merkel[6], and Casey S. Greene[1,2,7,8, ¶]

[1] Systems Pharmacology and Translational Therapeutics, University of Pennsylvania, Philadelphia, PA, USA
[2] Childhood Cancer Data Lab, Alex's Lemonade Stand Foundation, Philadelphia, PA, USA
[3] National Institute of Arthritis and Musculoskeletal and Skin Diseases, National Institutes of Health, Bethesda, MD, USA
[4] Division of Nephrology, Department of Internal Medicine, Michigan Medicine, Ann Arbor, MI, USA
[5] Department of Computational Medicine and Bioinformatics, Michigan Medicine, Ann Arbor, MI, USA
[6] Division of Rheumatology and the Department of Biostatistics, Epidemiology, and Informatics, University of Pennsylvania, Philadelphia, PA, USA
[7] Institute of Translational Medicine and Therapeutics, University of Pennsylvania, Philadelphia, PA, USA
[8] Institute of Biomedical Informatics, University of Pennsylvania, Philadelphia, PA, USA

[†] Current address: Childhood Cancer Data Lab, Alex's Lemonade Stand Foundation, Philadelphia, PA, USA
[¶] Correspondence to: greenescientist@gmail.com

SUMMARY

Unsupervised machine learning methods provide a promising means to analyze and interpret large datasets. However, most datasets generated by individual researchers remain too small to fully benefit from these methods. In the case of rare diseases, there may be too few cases available, even when multiple studies are combined. We sought to determine whether or not machine learning models could be constructed from large public data compendia and then transferred to small datasets for subsequent analysis. We trained models using Pathway Level Information ExtractoR (PLIER) over datasets of different types and scales. Models constructed from large public datasets were i) more detailed than those constructed from individual datasets; ii) included features that aligned well to important biological factors; iii) transferrable to rare disease datasets where the models describe biological processes related to disease severity more effectively than models trained within those datasets.

INTRODUCTION

The rapid expansion of the amount of publicly available gene expression data presents opportunities for discovery-driven research into rare diseases with poorly understood etiologies. Methods that extract a low-dimensional representation of the genome-scale gene expression data are useful for biological discovery, particularly in a multi-dataset setting. Such methods are biologically motivated, as genes that tend to vary together in their expression tend to function together. There are also technical considerations, as measuring combinations of genes will reduce the multiple hypothesis testing burden, can be useful for feature engineering, and is more likely to lead to more robust results than analyses of individual gene measurements (Cleary et al., 2017). These methods have advantages over two-group gene set-based comparisons because they provide more context for genes, are better fit to the underlying data, and remove the difficulty of identifying the most useful comparisons *a priori* (Stein-O'Brien et al., 2017a). Unsupervised machine learning (ML) methods including matrix factorization- and autoencoder-based approaches have successfully extracted biologically meaningful low-dimensional representations of gene expression data that can distinguish disease types, predict drug response, and identify new pathway regulators (Dincer et al., 2018; Stein-O'Brien et al., 2017b; Tan et al., 2017; Way and Greene, 2017).

Studying rare diseases using this data-intensive methodology is challenging. Even if multiple datasets exist for a disease of interest, the experimental designs may not be comparable or controls may not be selected by the same means. Combining, for instance, a clinical trial with longitudinal samples and no healthy controls and a study designed to test for differentially expressed genes between untreated cases and controls may not be appropriate. One strategy to integrate multiple, distinct datasets that cannot be directly combined is to identify modules or factors within each dataset and to seek to identify common changes within the set of identified factors. A challenge with this approach is that the factors, though they may capture similar processes, are not identical. Especially in the case where a one-to-one mapping for factors across datasets is lacking, this strategy can require human examination. This can be particularly time consuming because the number of comparisons grows with the square of the number of datasets.

Training a single, comprehensive model and using the same model for the analysis of datasets of interests allows the same factors to be directly compared. Moreover, signals may vary together in a particular small disease dataset but in fact be representative of distinct processes (e.g., glycolysis and an increase in infiltration of a specific immune cell type or change in tissue

composition). Models trained on these data alone may not distinguish correlated processes. However, with a large "shovel-ready" data source comprised of many conditions, including those where these processes are effectively "decoupled" (e.g., cell line experiments where glycolysis is perturbed but immune infiltration is irrelevant), it may be feasible to derive features from the large, heterogeneous compendium and transfer them to rare disease datasets. This is an example of *unsupervised transfer learning* and "feature-representation-transfer" (Pan and Yang, 2010). Intuitively, we might expect the features constructed from training on the data from heterogeneous conditions to be more generalizable and more appropriate in the case of disparate experimental designs because the model has trained on a broader set of contexts.

We wanted to determine if knowledge learned in a source domain—here, a large collection of human gene expression data—could be transferred to a target domain to improve the target task performance—here, the detection of patterns that characterize disease in unseen disease states. Specifically, we sought to extract a biologically meaningful low-dimensional representation of the recount2 compendium—a uniformly processed human RNA-seq dataset comprised of tens of thousands of samples and thousands of experiments (Collado-Torres et al., 2017)—and project antineutrophil cytoplasmic antibody (ANCA)-associated vasculitis (AAV) microarray datasets into this low-dimensional representation for downstream analysis. For knowledge transfer to be successful, the domains must be sufficiently related. We posited that transcriptional regulation, the resulting modular nature of gene expression data, and molecular mechanisms shared among diseases would support unsupervised transfer learning.

Just as the similarity between domains is important, selecting an appropriate method is crucial. We used the Pathway-Level Information ExtractoR (PLIER) matrix factorization framework (Mao et al., 2017). PLIER leverages prior biological knowledge (supplied as gene sets to the model), but can also identify features not associated with these sets. It learns patterns of correlated genes, termed latent variables (LVs). In this case the LVs are the low-dimensional representations. PLIER imposes certain constraints such that some of the learned LVs will align with the input gene sets. Because not *all* LVs are forced to relate to gene sets, some will capture major sources of variation due to technical factors. PLIER was shown to be robust to expression data processing choices, particularly for gene set-associated LVs (Mao et al., 2017), though it was not used in the context of a multi-dataset compendium. PLIER is well-suited for our task because capturing *shared biology* is key to transfer learning and because there will necessarily be technical noise in the recount2 compendium. We call this approach MultiPLIER because we train the model on multiple tissues and diseases. We demonstrate that *MultiPLIER* is an appropriate tool for biological discovery in multiple contexts, including cross-platform (microarray) applications.

To evaluate MultiPLIER, we designed experiments to assess the recovery of meaningful immune cell type-specific LVs, LVs associated with pathways that are perturbed during treatment with targeted therapeutics, and the concordance between MultiPLIER and disease- and dataset-specific models. We demonstrate that PLIER learns cell type-specific signatures and reduces technical batch effects when trained on a multi-dataset microarray compendium of systemic lupus erythematosus (SLE) whole blood (WB). We observe that MultiPLIER, trained on recount2, retains the biological fidelity of a compendium-specific model in the SLE WB evaluations and also recovers additional pathways. MultiPLIER is concordant with dataset-specific models in AAV and also identifies altered molecular processes in severe or active disease in multiple tissues. In summary, MultiPLIER's features have biological relevance and can be directly compared across datasets making it particularly valuable for the integrative analysis of complex human disease.

RESULTS

The MultiPLIER framework for unsupervised transfer learning

PLIER can be applied to or trained on individual datasets or large compendia (Mao et al., 2017). In Figure 1, we illustrate two different approaches to analyzing a multisystem disease with PLIER. PLIER automatically extracts LVs, shown as the matrix B, and their loadings (Z) (Figure 1A). We refer to the entire collection of LVs (B) as the latent space. In this example, three datasets from the same complex multisystem disease and distinct tissues are each used as a training set for a PLIER model. We refer to these PLIER models as _disease-specific or single dataset models_. This training approach results in three dataset-specific latent spaces (Figure 1A).

PLIER constrains the loadings Z to align with curated pathway or gene expression-derived cell type-specific gene sets, specified in the prior knowledge matrix C which is input to the model (Figure 1B). This ensures that some but not all of the resulting LVs capture _known biology_ and is an important and attractive feature of this particular approach. The prior information coefficient matrix U captures the relationships between LVs and the gene sets. In this example, the PLIER model is trained using the same input gene sets.

PLIER applies a penalty to U such that an individual LV should represent only a few gene sets (Figure 1C). The penalty makes LVs more interpretable. Ideally, an individual LV will be unambiguously associated with one pathway or cell type gene set or a collection of closely related gene sets. In Figure 1C, we show hypothetical LVs (columns) from the U matrix. In the top example, LV1 is strongly associated with lymphocyte, granulocyte, and cell cycle gene sets (Figure 1Ci). This association with many gene sets is undesirable because it hampers our ability to interpret the LV. Training on small datasets is more likely to result in latent variables that are associated with multiple pathways or cell types because there are fewer samples to disambiguate confounded factors. For instance, if neutrophil and T lymphocyte composition of all disease samples is increased relative to controls in a small dataset, these two signals may be difficult to distinguish. In the bottom example, LVs are each associated with a single pathway or cell type. In the bottom toy example, LVs are each associated with a single pathway or cell type—the ideal case (Figure 1Cii). PLIER additionally calculates AUC and p-value metrics for the association between a gene set and an LV.

In the case of multisystem disorders, identifying molecular processes that are perturbed across tissues is often desirable and may shed light on the pathobiology of that disease. In the dataset-specific PLIER case, we must map between individual dataset models—finding LVs that are similar to one another (Figure 1D). This can be challenging, as not all genes are represented in different datasets, multiple biological signals may be captured by a single LV if they tend to vary together (i.e., the neutrophil-T lymphocyte example above). Thus, an LV may or may not have a match across datasets, limiting the multi-tissue analyses that can be performed downstream.

An alternative to the disease-specific or single-dataset PLIER approaches is to train on a large collection of uniformly processed data from numerous biological contexts and conditions and then to use the loadings (Z) from this model to project the individual datasets into the _same latent space_ (Figure 1E). This is the approach we term _MultiPLIER_. In this work, we use the subset of the recount2 compendium (Collado-Torres et al., 2017) that includes metadata as our training set. This training set is comprised of approximately 37,000 samples, less than 100 of which are predicted to be from autoimmune disease and the majority of which are predicted to be from cell lines and tissue samples by MetaSRA (Bernstein et al., 2017).We refer to this model as either

MultiPLIER model or the full recount2 model to signify its training set and to LV1 from this model as MultiPLIER LV1. The large sample size and breadth of molecular processes represented increases the number of gene sets that the model captures and the *specificity* of the model—the learned LVs are more interpretable due to a "separation" of signals.

Once the individual datasets are in the same latent space, the LVs can serve as input into numerous downstream applications, including supervised tasks such as predicting response to treatment or testing for differential expression. In the Figure 1F toy example, MultiPLIER-learned LVs showed the same directionality (increased expression) in multiple tissues. This suggests that these molecular processes are systematically altered and warrants further investigation. For more information on implementation, see STAR Methods and the PLIER preprint (Mao et al., 2017).

A PLIER model trained on a systemic lupus erythematosus whole blood compendium learns SLE pathology-relevant LVs and differentiates biological signal from technical noise

We sought to determine if PLIER was appropriate for use with compendia assembled from multiple publicly available datasets. Gene expression studies from SLE have demonstrated the molecular heterogeneity of the disease (Banchereau et al., 2016) and studying WB allows us to evaluate how well the model captures immune cell-specific signals. We assembled an SLE WB compendium from 7 publicly available datasets that used different microarray platforms (Table 1). We selected these datasets because they included specific perturbations (e.g., treatments with a targeted biologic) or sample metadata (e.g., cell type counts) that supported evaluations of the model. In addition, some datasets had only submitter-processed data available, which can occur in repositories like Gene Expression Omnibus and ArrayExpress (Barrett et al., 2012; Edgar, 2002; Kolesnikov et al., 2014). We assembled the compendium by uniformly processing available raw data and scaling individual datasets (see STAR Methods for details). Despite this processing, dataset-specific factors were still evident in the first five principal components (PCs) (Figure S1).

We trained a PLIER model on this SLE WB compendium (termed the SLE WB PLIER model) and analyzed the learned LVs. Many cell types found in WB were represented (CD4 T cells in LV86, Monocytes in LV72 and LV109), as well as LVs associated with type I IFN signaling (LV6 and LV69) and highly relevant to autoimmunity like KEGG Graft vs. Host Disease and KEGG Antigen processing and presentation (LV59) (Figure 2A). Since PLIER does not require all learned LVs to be associated with gene sets, it can also learn major sources of variation associated with technical noise. As noted, the SLE WB compendium showed evidence of technical noise following processing (Figure S1).

We hypothesized that LVs associated with biological signals—those LVs with AUC > 0.75 for at least one gene set—would exhibit less technical noise. We performed principal component analysis (PCA) on different subsets of the latent space or B matrix for visualization purposes. With PCA on all latent variables (Figure 2B), we observe separation of the Agilent dataset (E-GEOD-78193) in PC1 and PC2. Once the B matrix was subset to only LVs associated with pathways or cell types, less separation between datasets was observed (Figure 2C). Note that E-GEOD-65391 is a pediatric cohort, which may account for some of the separation in PC2 in Figure 2C. In contrast, when examining only LVs that are not explicitly tied to gene sets (Figure 2D), the dataset-specific effect is more exaggerated than in all LVs (Figure 2B). This suggests that this model partitions the biological signal from variables that capture technical noise and potentially biological signal. Furthermore, the model learns LVs associated with pathways that are highly relevant to SLE, even when trained on a compendium comprised of multiple datasets on different platforms.

As illustrated in Figure 2A, PLIER learned cell type-specific signals from heterogeneous expression data. PLIER was also shown to outperform other methods (e.g., CIBERSORT, Newman et al., 2015; non-negative matrix factorization) in estimating cell type proportions from a small dataset of whole blood RNA-seq samples as measured by correlation to counts from mass cytometry data (Mao et al., 2017). We sought to evaluate whether or not these signals captured in the SLE WB mode were well-correlated with cell type counts, despite the fact that the model was not explicitly trained to predict cell type composition. We used the Banchereau, et al. dataset (E-GEOD-65391) from the SLE WB compendium which includes neutrophil counts in the sample metadata for this evaluation. We selected LV87 from the SLE WB PLIER model (SLE WB LV87) from multiple LVs that were significantly associated (FDR < 0.05) with neutrophil gene sets because it lacked an association with other myeloid cell types (e.g., monocyte- or macrophage-related gene sets) and could then be considered to be specific to neutrophil gene sets. We observed that SLE WB LV87 showed agreement (R^2 = 0.29) with the E-GEOD-65391 neutrophil counts (Figure 3A).

MultiPLIER learns cell type-associated LVs that are well-correlated with counts or estimates in multiple tissues and disease contexts

Systemic lupus erythematosus is not a rare disorder. Studying SLE does not come with the same challenges as studying a rare disorder with a data intensive approach—namely, the small sample size and lack of publicly available data. This is evidenced by the fact that we were able to assemble a compendium comprised of 1640 WB samples without selecting *all* publicly available data. We designed MultiPLIER, which is constructed from a broad compendium of human RNA-seq data, to study complex rare diseases in particular. To compare SLE WB PLIER and MultiPLIER, we projected the SLE WB compendium into the MultiPLIER latent space. We evaluated neutrophil counts using MultiPLIER LV603, which was significantly associated (FDR < 0.05) with several neutrophil gene sets and not significantly associated with gene sets from other myeloid lineage cell types (Figure S2A). We observed stronger agreement between MultiPLIER LV603 expression and neutrophil counts (R^2 = 0.36) than between SLE WB LV87 and the counts.

This agreement with the Banchereau, et al. data was not limited to the neutrophil signatures (Banchereau et al., 2016). In the original publication, the authors stratified patients into three groups based on their SLEDAI (SLE Disease Activity Index) scores (Rovenský and Payer, 2009) and demonstrated that those patients with higher disease activity had increase plasmablast counts (as measured by FACS in a subset of patients). Both the SLE WB PLIER and MultiPLIER plasma cell associated LVs showed the same pattern as was reported in the original publication (Figure S2B-C) (Banchereau et al., 2016). MultiPLIER recovers these features even though it was not trained on this dataset and has not encountered the particular technical noise in the SLE WB compendium. These results also strongly suggest that MultiPLIER is appropriate for cross-platform analyses, as recount2 is RNA-seq data.

To assess whether or not this performance was generalizable to other datasets or disease conditions (including solid tissues), we performed additional evaluations in a nasal brushing microarray dataset (NARES dataset) that included patients with granulomatosis with polyangiitis (GPA), sarcoidosis, and healthy controls (Grayson et al., 2015) and a microarray dataset of isolated immune cell types from different autoimmune conditions (E-MTAB-2452; McKinney et al., 2015). Neutrophil counts were unavailable for the NARES solid tissue dataset. Instead, we use MCPcounter (Becht et al., 2016), a state-of-the-art method for estimating immune cell infiltration from solid tissues including non-cancerous tissues, as our dependent variable. MultiPLIER LV603

is highly correlated (R^2 = 0.9) with the MCPcounter neutrophil estimate in the NARES dataset (Figure 3C). The isolated leukocyte dataset include samples that were selected for CD4 (T cells), CD14, (monocytes) and CD16 (neutrophils). MultiPLIER LVs associated with particular cell type gene sets were generally most highly expressed in the samples comprised of that isolated cell type (Figure S3). Taken together, this demonstrates that MultiPLIER captures highly relevant cell type-specific information in microarray datasets.

MultiPLIER learns type I and type II interferon latent variables that are perturbed in a manner consistent with interferon blockade in SLE clinical trials

We have established that MultiPLIER captures cell type-specific signals. However, the gene sets we used with PLIER were not limited to cell type-specific gene sets and the SLE WB PLIER model learned LVs associated with pathways such as KEGG Interferon Alpha/Beta signaling (Figure 2A). The SLE WB compendium includes two clinical trials of targeted treatments that block interferon (IFN) signaling. One, IFN-alpha-kinoid (IFN-K), is a therapeutic vaccine that results in polyclonal antibodies to IFN-alpha—a type I IFN (E-GEOD-39088; Lauwerys et al., 2013). The other trial is of AMG811, a monoclonal antibody to IFN-gamma—the type II IFN (E-GEOD-78193; Welcher et al., 2015). The expression patterns induced by different types of IFN can be difficult to distinguish (Chiche et al., 2014). Here, we assessed MultiPLIER's ability to distinguish the two types of IFN signaling.

MultiPLIER learned two LVs, LV116 and LV140, that were strongly associated with type I (IFN-alpha/beta) and type II (IFN-gamma) signaling pathways, respectively (Figure S2A). We projected expression data from Lauwerys, et al. (Lauwerys et al., 2013), a clinical trial of the IFN-alpha-kinoid (IFN-K) therapeutic that results in the blockade of type I IFN via polyclonal antibody response, into the MultiPLIER latent space. The authors of the original IFN-K publication stratified patients with SLE into two groups: IFN-negative (n = 9), who had IFN-alpha induced gene expression signatures that resembled healthy controls, and IFN-positive patients (n = 18), who had higher expression of the IFN-alpha signature they derived through *ex vivo* stimulation of healthy control blood with IFN-alpha. (Group labels from the original publication were not associated with the accession.) The type I IFN LV116 separates patients into two groups (Figure 4A); this separation is not evident in type II IFN LV140 (Figure 4A). The SLE WB PLIER model did not learn an LV significantly associated with type II IFN signaling (e.g., REACTOME Interferon gamma signaling).

Treatment of IFN-positive patients with IFN-K showed a significant decrease in LV116 (Type I IFN) expression at day 112 as compared to the placebo group (p = 0.017; Figure 4B). This is consistent with results using the *whole blood-specific* modular transcriptional repertoire analysis approach popular in rheumatology studies of WB expression data (Chaussabel et al., 2008; Chiche et al., 2014; Oswald et al., 2015) (Figure S4A-B). MultiPLIER also stratifies patients into IFN-positive and IFN-negative groups using the SLE WB PLIER type I IFN LVs expression (LV6, LV69, LV110; Figure S4C), but the LV with the highest AUC for REACTOME Interferon alpha/beta signaling—LV6—did not show a reduction in expression in IFN-positive patients at day 112 (Figure S4D), though LV110 showed a significant difference (p = 0.022).

To address type II IFN signaling response to therapy, AMG811 response data, including *ex vivo* stimulation of healthy blood with IFN-gamma, included in the SLE WB compendium was interrogated. The signatures (e.g., LVs or modules) that captured type II IFN signaling showed a trend of decreased expression during treatment with AMG811, whereas the type I IFN signatures were unchanged (Figure S5). In summary, MultiPLIER distinguishes type I and type II IFN

signatures that are not detected in a PLIER model trained only on SLE WB data and MultiPLIER does not have the tissue-specificity and lack of information about individual genes' contributions limitations that are inherent to the modular transcriptional frameworks.

Subsampling of the recount2 compendium reveals that both sample size and breadth of biological conditions contribute to MultiPLIER's characteristics

We wanted to determine whether the size of the recount2 compendium or the broad sampling of human molecular processes contained within contributed to the characteristics of the MultiPLIER model. To assess this, we randomly selected samples from the recount2 compendium such that it was the same size of the SLE WB compendium (n = 1640). We repeated this procedure 10 times and trained a PLIER model on each subset. We also trained models on the entire SLE WB compendium initialized with 10 different random seeds. We evaluated models by the following metrics: the number of LVs (k), the proportion of input gene sets that have at least one LV significantly associated (FDR < 0.05) with them (termed "pathway coverage"), and the proportion of LVs that are significantly associated with at least one pathway (FDR < 0.05). The results are shown in Figure 5.

Models trained on the SLE WB compendium had a larger k (Figure 5A); k is set by first determining the number of significant PCs (or major variance components) identified in the dataset. This result means that there are fewer statistically significant PCs in the subsampled recount2 sets, which may be due to the heterogeneous tissue composition of the SLE WB samples. In addition, k is more variable in the subsampling case. This is as expected because the training set is comprised of 1640 random samples. Also as expected, the full recount2 model had larger k than any dataset of 1640 samples (Figure 5A). The proportion of pathways or gene sets that are significantly associated with LVs is ~0.2 for all the models under consideration and is much lower than the pathway coverage of the full recount2 model (Figure 5B), suggesting that coverage was constrained by sample size. Despite there being fewer LVs in the subsampled models, a higher proportion of the LVs were significantly associated (FDR < 0.05) with pathways or gene sets (Figure 5C) as compared to the SLE WB models. These results suggest that including a broader set of biological conditions in a training set may result in more detectable pathway variation. Interestingly, the full recount2 PLIER model has a lower proportion of LVs significantly associated with prior knowledge. This could mean that the full recount2 PLIER model captures novel biology, has more technical noise introduced (as samples are from many different experiments), or both.

The latent space of the MultiPLIER model shows agreement with a PLIER model trained on a small vasculitis cohort and captures an autoantigen signature

The MultiPLIER approach is most attractive for rare disease where sample sizes are limited. AAV is a group of rare diseases characterized by neutrophil-mediated inflammation of small and medium vessels in multiple systems, and includes the sub-types of granulomatosis with polyangiitis (GPA), microscopic polyangiitis (MPA), and eosinophilic granulomatosis with polyangiitis (EGPA; formerly Churg-Strauss). We projected microarray datasets from multiple tissues in AAV—nasal brushings, microdissected glomeruli from kidney, and blood—into the MultiPLIER latent space and examined the expression of the MultiPLIER LVs in these data. To our knowledge, the recount2 compendium contains no vasculitis data.

We evaluated the extent to which the MultiPLIER model, trained without any vasculitis data, captured vasculitis-relevant features. We trained a PLIER model on the NARES nasal brushing

dataset (n = 79; Grayson et al., 2015), and compared the results to the projection into the MultiPLIER latent space. NARES includes data from patients with GPA, a form of AAV. First, we identified the MultiPLIER LV that was the best match for each of the 34 LVs learned by the NARES PLIER as determined by correlation between the loadings (highest Pearson correlation coefficient). We observed that the expression levels of the best matched LVs were positively correlated between the NARES and MultiPLIER latent spaces (i.e., B matrices) and much more so than we would expect than any random combination of LVs between models (Figure 6). Figure S6 displays the correlation between LV expression alongside the correlation between loadings for the best match LVs. Those NARES LVs that are significantly associated (FDR < 0.05) with at least one gene set have particularly strong agreement (blue points in Figure 6). This supports the hypothesis that biological features captured by MultiPLIER are generalizable and appropriate for transfer learning.

We identified differentially expressed MultiPLIER LVs (DELVs) in the peripheral blood mononuclear cell (PBMC) fraction from Cheadle, et al. (Cheadle et al., 2010). This dataset included patients with GPA (known as Wegener's granulomatosis at the time) and healthy controls. Patients with GPA were stratified into two groups in the original publication: those with expression patterns distinct from healthy control (here, termed "GPA-positive") and those that were more similar to healthy controls (termed "GPA-negative"). We tested for differential expression between these three groups—GPA-positive, GPA-negative, and control—and identified MultiPLIER LV599 as differentially expressed (FDR = 1.17e-08; Figure S7A). ANCA antigens (i.e., the neutrophilic granule components that autoantibodies react with) are among the genes that contribute most to MultiPLIER LV599 (black bars in Figure S7D). This demonstrates the fact that MultiPLIER captures biology shared in multiple tissue contexts (i.e., granule development) that is aberrant in GPA without being trained on GPA data.

A MultiPLIER meta-analysis framework identifies differentially expressed in three tissues affected by ANCA-associated vasculitis

We were interested in molecular processes perturbed in the systemic disease AAV. Because we projected the three tissue datasets into the same latent space, MultiPLIER lends itself naturally to the meta-analysis of the NARES, glomeruli, and blood data as we did not have to attempt to match LVs across models making analysis much less time consuming for multiple datasets. PLIER reduces technical noise (Figure 2C) and discards genes and gene sets that are irrelevant to the training data (Mao et al., 2017), which are distinct advantages over other gene set methods such as Gene Set Enrichment Analysis (GSEA; Subramanian et al., 2005). We performed differential expression analysis to identify DELVs in each of the three datasets and to find those LVs that had similar expression patterns across tissues.

We examined signatures of more severe or active disease in AAV across organ systems. In the blood dataset, patients with a GPA-positive signature (most dissimilar from healthy controls) had higher Birmingham Vasculitis Activity Scores (BVAS; Cheadle et al., 2010). In the NARES dataset, there were three groups of patients with GPA—those with *active* nasal disease, those with *prior* nasal disease, and those that had *no history* of nasal disease—and a composite comparator group of patients with eosinophilic granulomatosis with polyangiitis (EGPA, Churg-Strauss), allergic rhinitis, sarcoidosis, and healthy controls (Grayson et al., 2015). The glomeruli dataset contained patients with AAV as well as those with nephrotic syndrome and living donor controls (Grayson et al., 2018). We used all available groups for comparisons (see STAR Methods). We identified 22 MultiPLIER LVs that are differentially expressed (FDR < 0.05) in all three datasets (Figure S7A) and additional LVs differentially expressed in two out of three tissues.

We were particularly interested in LVs that showed evidence of increased expression in GPA with active disease (NARES/nasal brushings), the GPA-positive group (PBMCs), and in the vasculitis group in kidney. MultiPLIER LV10, which is significantly associated with the M0 macrophage gene set from Newman, et al. (Newman et al., 2015), was overexpressed in these three groups relative to controls (Figure 7A-C; multigroup comparison FDR < 0.05; p-value from Wilcoxon rank sum test shown in figure). Top (highest weight) genes for MultiPLIER LV10 (Figure 7B) include the M2 macrophage marker *CD163* and *CCL2* which has been shown to influence macrophage polarization to an M2 phenotype (Sierra-Filardi et al., 2014). The top 25 genes also include genes related to tissue remodeling such as *MMP9* and *MMP14* (Figure 7D). MultiPLIER LV937 was significantly associated with a HIF-1a transcription factor network and showed a differential expression similar to MultiPLIER LV10 pattern (Figure S8B-D). We previously demonstrated a signature associated with aberrant metabolism is present in AAV glomeruli and identified macrophages as the most likely cellular source (Grayson et al., 2018). These results suggest that this immunometabolism process may be occurring in multiple tissues in AAV and establishes the benefit of meta-analyses with MultiPLIER.

DISCUSSION

We describe MultiPLIER—an unsupervised transfer learning framework for identifying perturbed molecular processes in complex human disease across organ systems. MultiPLIER models are generated by applying a robust factorization approach named PLIER (Mao et al., 2017) to a large, uniformly processed body of gene expression from multiple tissues and biological processes (Collado-Torres et al., 2017). MultiPLIER captured cell type-specific patterns as well as a model trained on the SLE whole blood compendium, despite the inclusion of very few SLE WB samples in recount2. It learned an LV associated with Type II IFN signaling; the SLE WB PLIER model did not. MultiPLIER analysis of multiple tissues from AAV allowed revealed signatures indicative of altered immunometabolism. Thus, we have demonstrated that training a model on data that represents an arbitrary sampling of human physiological processes captures well-documented features similarly to a tissue- and disease-specific model in a common disease, and learns additional biological factors or pathways, which aids in the interpretation of rare disease expression data from multiple tissues.

MultiPLIER learned LVs that were significantly associated with macrophage and HIF-1alpha TF network gene sets. These LVs had increased expression in more severe or active AAV in multiple tissues. The microdissected glomeruli dataset included in these analyses served as the discovery cohort in Grayson, et al. (Grayson et al., 2018) where we demonstrated altered metabolic pathway expression was tightly coupled with myeloid lineage transcriptional patterns. In that work, we observed colocalization of HIF-1alpha with the tissue macrophage marker CD68 in AAV kidney biopsies. These results suggest that the alteration of metabolism in tissues in patients with AAV should be explored, though they may be affected by treatment with glucocorticoids. Although we highlight two MultiPLIER LVs that show increased expression in multiple tissues from AAV, we expect that there are additional LVs of interest that warrant investigation including an LV associated with a neutrophil gene set (MultiPLIER LV524), a cell type highly relevant to these diseases.

PLIER leverages the underlying modular structure of gene expression data. MultiPLIER expands on this by taking advantage of the biological processes shared among different conditions. We expect that the benefits of training on a large compendium of uniformly processed data—the first

of which is the reduction of processing time associated with assembling a disease-specific compendium even when feasible—are not limited to PLIER. We speculate that other matrix factorization approaches and autoencoders that have been shown to extract a biologically meaningful low dimensional representation from gene expression data (Dincer et al., 2018; Stein-O'Brien et al., 2017b; Tan et al., 2017; Way and Greene, 2017) are appropriate for a similar transfer learning approach. The constraints that force the PLIER model to learn biological signal made it an attractive method for this work. Indeed, the LVs learned by MultiPLIER that were significantly associated with gene sets were shown to be in strong agreement with dataset-specific model LVs. Future investigation into a variety models with different strengths (e.g., those that capture nonlinear patterns), with particular emphasis on the factors in the latent space that are well-aligned with biological signatures of interest, is supported by this work.

There are several limitations of this study. For instance, we made minimal effort to adjust the data distributions between platforms (i.e., RNA-seq and microarray) (Tan et al., 2017; 2016; Thompson et al., 2016) or impute the expression of genes missing between models. Despite these areas for improvement, studying the expression of MultiPLIER LVs in microarray datasets from rare disease proved to be useful. We also presented the results for one MultiPLIER model trained on the recount2 compendium and PLIER is non-deterministic. An ensemble approach, like that introduced in Tan, et al. (Tan et al., 2017) could be beneficial, though our repeats of SLE WB PLIER training suggest that PLIER produced stable models at least in terms of the number of latent variables and proportion of input pathways captured. Additionally, any approach that transfers patterns learned in a large compendium to an "unseen" disease dataset runs the risk of missing subtle disease-specific and disease-affected tissue-specific signals. Conducting hypothesis-generating research in small datasets in general poses a risk for false discovery. After first comparing model features to well-characterized datasets to validate model construction, we focus on patterns that are present in multiple datasets and tissues.

The union of automatic pattern extraction using unsupervised machine learning with the exploitation of shared biological processes opens up several exciting avenues for research in complex human diseases and specifically in the rare or understudied disease space. As large compendia of uniformly processed gene expression data become more common, cross-organism transfer learning could prove fruitful for identifying appropriate models for the study of human disease and detecting targeted perturbations reminiscent of human pathophysiology. Furthermore, we hypothesize that the latent variables learned by MultiPLIER may be useful and generalizable input features for supervised machine learning problems (Dincer et al., 2018) including predicting response to treatment in small clinical trials. MultiPLIER is a powerful transfer learning strategy for studying systems and discovery-driven research into the mechanisms underlying human disease.

Although we propose the MultiPLIER is a particularly strong tool for the study of rare diseases, in which the size of gene expression datasets are inherently small compared to similar data derived from common diseases, we also foresee that MultiPLIER might be quite useful for the increasing study of precision (personalized) medicine and the now routine recognition of newly-discovered single-gene subsets of diseases. The precision medicine approach is linked to the growing appreciation of the importance of identifying subsets of patients within larger populations of patients with a specific disease. Such subsets may identify patients with i) different sets of clinical manifestations of disease; ii) varying degrees of disease severity; iii) different long-term outcomes; and iv) varied responses to specific therapies. Such disease subsets almost certainly have some differences in genetic background and/or gene expression in involved tissues. Thus, these approaches to precision medicine will essentially transform some common diseases into

subsets of rarer diseases, many of which could benefit from the MultiPLIER approach to data analysis.

METHODS

Gene expression data

recount2 compendium

We obtained gene-level RangedSummarizedExperiment data from the recount2 compendium (Collado-Torres et al., 2017) using the recount Bioconductor package and calculated Reads Per Kilobase Million (RPKM). We excluded samples without metadata; a total of 37,032 samples were included. Individual experiments were concatenated without any additional processing and genes were z-scored for use with PLIER. We converted from Ensembl gene IDs to gene symbols using the biomaRt Bioconductor package.

To describe the subset of recount2 data that was used to train the model, we used MetaSRA database (Bernstein et al., 2017) (v1.4). MetaSRA uses a custom pipeline to predict sample source (e.g., cell line) and label with ontology terms from sources such as the Experimental Factor Ontology (Malone et al., 2010). We mapped run accessions (utilized in recount2) to sample accessions (utilized in MetaSRA) with the SRAdb (Zhu et al., 2013). Approximately 81% of samples in the training compendium were also present in MetaSRA.

Microarray data

The NARES nasal brushing dataset, the isolated immune cell data (E-MTAB-2452) and the following systemic lupus erythematosus datasets were normalized using Single Channel Array Normalization (SCANfast implementation) and Brainarray v22.0.0 (Entrez ID version; Dai, 2005): E-GEOD-39088, E-GEOD-61635, E-GEOD-72747, and E-GEOD-11907. The NARES dataset was adjusted for batch effect with ComBat (Johnson et al., 2007). The glomeruli dataset was annotated using Brainarray v19.0.0, quantile normalized, and batch corrected using ComBat (Grayson et al., 2018) . The publicly available submitter-processed data was used for all other datasets (e.g., SLE WB [E-GEOD-49454, E-GEOD-65391, E-GEOD-78193], GPA PBMCs [GSE18885]). The GPA PBMCs data was z-scored for use with PLIER and MultiPLIER. All conversion from Entrez IDs to Gene symbols (the input to PLIER) was performed using the org.Hs.eg.db Bioconductor package.

Systemic lupus erythematosus whole blood compendium

The Illumina and Agilent submitter-processed data gene identifiers were converted to Entrez IDs using the illuminaHumanv4.db and hgu4112a.db packages, respectively. Duplicate Entrez IDs were collapsed to the mean expression value. Paired samples from E-GEOD-11907 (Affymetrix HGU133A and HGU133B) were concatenated. All SCAN-processed Affymetrix data was concatenated and the submitter-processed datasets were normalized using the Affymetrix target quantiles (preProcessCore Bioconductor package). Genes were $[0,1]$ scaled in each dataset before joining; only genes present in each dataset were retained.

Training Pathway-Level Information Extractor (PLIER) models and evaluation

PLIER training

We used the PLIER R package (commit: a2d4a2) to train PLIER models. The number of latent variables (k parameter) was set by determining the number of significant principal component using the PLIER num.pcs function and, following the authors recommendations, including an additional 30%. The expression matrix used as input to the PLIER function was z-scored using the PLIER rowNorm function. We used the bloodCellMarkersIRISDMAP, svmMarkers, and canonicalPathways data included in the package as the prior information. We used the default settings of the PLIER function, which automatically sets the L1 and L2 parameters (see Mao et al., 2017 for more information about implementation).

Subsampling experiment

For the subsampling experiments (Figure 5), we randomly selected 1640 samples from the recount2 dataset as our new expression matrix and treated it as described above. We trained 10 PLIER models on the subsampled recount2 dataset (n = 1640) and the full SLE WB compendium (n = 1640) initialized with different random seeds. Training was performed as described above. To evaluate the models, we calculated the proportion of input pathways that were significantly associated (FDR < 0.05) with at least one LV—termed "pathway coverage"—and the proportion of the total LVs that were significantly associated (FDR < 0.05) with at least one gene set.

MultiPLIER

We summarize the relevant details from the PLIER method below. For more information about implementation, see Mao, et al. (Mao et al., 2017). Given the gene expression matrix Y and the prior information matrix (e.g., gene sets) C, PLIER finds the prior information coefficient matrix U, the loadings matrix Z, and the latent variables matrix B, minimizing:

$$||Y - ZB||^2_F + \lambda_1||Z - CU||^2_F + \lambda_2||B||^2_F + \lambda_3||U||_{L^1}$$

where $U > 0$ and $Z > 0$. The terms are minimizing the reconstruction error, enforcing that Z represents sparse combinations of the input gene sets, an L2 penalty on B and an L1 penalty on U, respectively. As noted above, we used the default settings for λ_1 and λ_2 for a given training set when training a PLIER model (see Mao et al., 2017 for more information).

The final B matrix returned by PLIER is calculated as follows:

$$B = (Z^T Z + \lambda_2 I)^{-1} Z^T Y$$

We use the Z matrix and the λ_2 parameter from the source domain model to calculate the expression values of the target samples in the source domain latent space B_{target} below:

$$B_{target} = (Z^T Z + \lambda_2 I)^{-1} Z^T Y_{target}$$

Where Y_{target} is the row-normalized (z-scored) expression matrix from the dataset (target domain) that we would like to project into the latent space learned by the source domain model. If genes present in the training set (source domain) were absent from the target gene expression matrix, we set the expression values for the missing genes to zero for all samples. As the expression data is z-scored, this is in essence setting the expression values to the mean. We calculated

B_{target} for the following microarray datasets: SLE WB compendium, NARES nasal brushing, microdissected glomeruli (kidney), GPA peripheral blood mononuclear cells, isolated leukocytes (CD4, CD14, CD16 selection).

Latent variable differential expression analysis

We performed differential expression analysis using the limma Bioconductor package. We performed contrasts between all groups and used Benjamini-Hochberg adjustment for multiple hypothesis testing ("BH" method). We considered all LVs with FDR < 0.05 to be differentially expressed.

Principal components analysis

We performed principal components analysis (PCA) on the fully processed SLE WB compendium (Figure S1), as well as subsets of the latent space (B matrix) from the SLE WB PLIER model using the prcomp function in R. The subsets of B were as follows: all LVs ("All Latent Variables" Figure 2B), only LVs with AUC > 0.75 for at least one gene sets ("Pathway-associated Latent Variables" Figure 2C), only LVs with AUC $\leq$ 0.75 for all gene sets (Figure 2D).

Cell type evaluations

Neutrophil-associated LVs were identified based on the AUC and FDR (< 0.05) values with neutrophil gene sets. We selected neutrophil-associated LVs for display in Figure 3 based on a lack of significant association with monocyte-, macrophage-, or granulocyte progenitor-related gene sets, as this suggests specificity for neutrophils rather than a general myeloid lineage association. We used the neutrophil counts from the metadata associated with the publicly available accession (E-GEOD-65391) for Banchereau, et al. (Banchereau et al., 2016). We calculated the MCPcounter neutrophil estimate with the MCPcounter.estimate function in the MCPcounter R package (Becht et al., 2016). We used the R linear model function (lm) to calculate the R^2 values in the neutrophil scatterplots in Figure 3. Disease activity group labels (Figure S2) were also obtained from the E-GEOD-65391 sample metadata. Plasma cell-associated LVs were identified by their significant (FDR < 0.05) association with plasma cell gene sets; LVs with the highest AUC were selected for display. Pairwise t-tests between disease activity groups were performed with the pairwise.t.test function in R. For the heatmap in Figure S3, we selected LVs for display if they had one of the following strings in their names: CD4, monocyte, or neutrophil.

Interferon clinical trials

Whole blood modular framework

We used the IFN-related modules reported in Chiche, et al. (Chaussabel et al., 2008; Chiche et al., 2014; Obermoser et al., 2013). Briefly, in this approach, modules are constructed by clustering gene expression data from whole blood from patients with a variety of autoimmune or infectious conditions and identifying shared expression patterns (Chaussabel et al., 2008). The result is sets of genes with correlated expression or modules. It is important to note that these modules are then somewhat specific to whole blood. We obtained the gene sets from http://www.biir.net/public_wikis/module_annotation/V2_Trial_8_Modules and converted to Entrez ID using Tribe (Zelaya et al., 2016). (These converted gene sets are available in our analysis Github repository.) We summarized the expression of a module by taking the mean expression value for genes in the module for a sample.

PLIER models

We selected IFN-related LVs based on their association with the following genesets: REACTOME Interferon signaling, REACTOME Interferon alpha/beta signaling (Type I), REACTOME Interferon gamma signaling (Type II) (Fabregat et al., 2017; Subramanian et al., 2005).

IFN-K trial analyses

In our analyses of the IFN-alpha-kinoid trial (E-GEOD-39088; Lauwerys et al., 2013), we stratified the patients into IFN-positive and IFN-negative groups to be consistent with the original publication and therefore facilitate our evaluations. We did not have the labels from the original publication. We identified IFN-negative samples based on baseline expression values and using the following criteria: 9 patients with the lowest summary expression values for M1.2 (whole blood modular framework), 9 patients with the lowest expression of MultiPLIER LV116 (MultiPLIER), or 9 patients that had the lowest expression in at least 2 of the SLE WB LVs—LV6, LV69, LV110 (SLE WB PLIER). All other patients were considered IFN-positive. Patients that received placebo were placed in a "placebo" group for analysis. We calculated the change in expression value between baseline and days 112 and 168 for each patient.

Latent space agreement

To determine whether or not the NARES PLIER model and MultiPLIER latent space expression values were in agreement, we first identified the best match LVs between the two models. For each pair of LVs between the two models, we calculated the Pearson correlation coefficient between the loadings (Z). The MultiPLIER LV with the highest correlation value to an individual LV from the NARES model was considered the best match. We calculated the Pearson correlation values between the B matrix learned by the NARES PLIER model and B from the projection of the NARES data into the MultiPLIER latent space.

ACKNOWLEDGMENTS

We would like to thank Gregory P. Way for code review and thoughtful discussion.

DATA AND SOFTWARE AVAILABILITY

Code for expression data processing and MultiPLIER training is available at https://github.com/greenelab/rheum-plier-data. Code for analysis is available at https://github.com/greenelab/multi-plier. MultiPLIER training data and model are available at figshare with the DOI: 10.6084/m9.figshare.5716033.v4; A release (v0.1.0) of the data, with the exception of the MultiPLIER model and training data, and code is available at DOI: 10.6084/m9.figshare.6982919. We have prepared a Docker image with the software dependencies required to reproduce the analysis: https://hub.docker.com/r/jtaroni/multi-plier/ (tag 0.1.0).

FUNDING

This work was supported in part by grants to CSG from the Gordon and Betty Moore Foundation (GBMF 4552), the National Institutes of Health (R01 HG010067), the Chan Zuckerberg Initiative Donor Advised Fund of the Silicon Valley Community Foundation (2018-182718), and Alex's Lemonade Stand Foundation (Childhood Cancer Data Lab). This work was supported in part by

The Vasculitis Clinical Research Consortium (VCRC) (U54 AR057319 and R01 AR064153) which is part of the Rare Diseases Clinical Research Network (RDCRN), an initiative of the Office of Rare Diseases Research (ORDR), National Center for Advancing Translational Science (NCATS). The VCRC is funded through collaboration between NCATS, and the National Institute of Arthritis and Musculoskeletal and Skin Diseases, and has received funding from the National Center for Research Resources (U54 RR019497). JNT received support from T32 AR007442.

REFERENCES

Abbas, A.R., Baldwin, D., Ma, Y., Ouyang, W., Gurney, A., Martin, F., Fong, S., van Lookeren Campagne, M., Godowski, P., Williams, P.M., 2005. Immune response in silico (IRIS): immune-specific genes identified from a compendium of microarray expression data. Genes and immunity 6, 319–331.

Banchereau, R., Hong, S., Cantarel, B., Baldwin, N., Baisch, J., Edens, M., Cepika, A.-M., Acs, P., Turner, J., Anguiano, E., Vinod, P., Kahn, S., Obermoser, G., Blankenship, D., Wakeland, E., Nassi, L., Gotte, A., Punaro, M., Liu, Y.-J., Banchereau, J., Rossello-Urgell, J., Wright, T., Pascual, V., 2016. Personalized Immunomonitoring Uncovers Molecular Networks that Stratify Lupus Patients. Cell 165, 551–565. doi:10.1016/j.cell.2016.03.008

Barrett, T., Wilhite, S.E., Ledoux, P., Evangelista, C., Kim, I.F., Tomashevsky, M., Marshall, K.A., Phillippy, K.H., Sherman, P.M., Holko, M., Yefanov, A., Lee, H., Zhang, N., Robertson, C.L., Serova, N., Davis, S., Soboleva, A., 2012. NCBI GEO: archive for functional genomics data sets—update. Nucleic Acids Research 41, D991–D995. doi:10.1093/nar/gks1193

Becht, E., Giraldo, N.A., Lacroix, L., Buttard, B., Elarouci, N., Petitprez, F., Selves, J., Laurent-Puig, P., Sautès-Fridman, C., Fridman, W.H., de Reyniès, A., 2016. Estimating the population abundance of tissue-infiltrating immune and stromal cell populations using gene expression. Genome Biol. 17, 218. doi:10.1186/s13059-016-1070-5

Bernstein, M.N., Doan, A., Dewey, C.N., 2017. MetaSRA: normalized human sample-specific metadata for the Sequence Read Archive. Bioinformatics 33, 2914–2923. doi:10.1093/bioinformatics/btx334

Chaussabel, D., Quinn, C., Shen, J., Patel, P., Glaser, C., Baldwin, N., Stichweh, D., Blankenship, D., Li, L., Munagala, I., Bennett, L., Allantaz, F., Mejias, A., Ardura, M., Kaizer, E., Monnet, L., Allman, W., Randall, H., Johnson, D., Lanier, A., Punaro, M., Wittkowski, K.M., White, P., Fay, J., Klintmalm, G., Ramilo, O., Palucka, A.K., Banchereau, J., Pascual, V., 2008. A modular analysis framework for blood genomics studies: application to systemic lupus erythematosus. Immunity 29, 150–164. doi:10.1016/j.immuni.2008.05.012

Cheadle, C., Berger, A.E., Andrade, F., James, R., Johnson, K., Watkins, T., Park, J.K., Chen, Y.-C., Ehrlich, E., Mullins, M., Chrest, F., Barnes, K.C., Levine, S.M., 2010. Transcription of proteinase 3 and related myelopoiesis genes in peripheral blood mononuclear cells of patients with active Wegener's granulomatosis. Arthritis & Rheumatism 62, 1744–1754. doi:10.1002/art.27398

Chiche, L., Jourde-Chiche, N., Whalen, E., Presnell, S., Gersuk, V., Dang, K., Anguiano, E., Quinn, C., Burtey, S., Berland, Y., Kaplanski, G., Harle, J.-R., Pascual, V., Chaussabel, D., 2014. Modular transcriptional repertoire analyses of adults with systemic lupus erythematosus reveal distinct type I and type II interferon signatures. Arthritis Rheumatol 66, 1583–1595. doi:10.1002/art.38628

Cleary, B., Cong, L., Cheung, A., Lander, E.S., Regev, A., 2017. Efficient Generation of Transcriptomic Profiles by Random Composite Measurements. Cell 171, 1424–1436.e18. doi:10.1016/j.cell.2017.10.023

Collado-Torres, L., Nellore, A., Kammers, K., Ellis, S.E., Taub, M.A., Hansen, K.D., Jaffe, A.E., Langmead, B., Leek, J.T., 2017. Reproducible RNA-seq analysis using recount2. Nat Biotechnol 35, 319–321. doi:10.1038/nbt.3838

Dai, M., 2005. Evolving gene/transcript definitions significantly alter the interpretation of GeneChip data. Nucleic Acids Research 33, e175–e175. doi:10.1093/nar/gni179

Dincer, A.B., Celik, S., Hiranuma, N., Lee, S.-I., 2018. DeepProfile: Deep learning of cancer molecular profiles for precision medicine. bioRxiv 278739. doi:10.1101/278739

Ducreux, J., Houssiau, F.A., Vandepapelière, P., Jorgensen, C., Lazaro, E., Spertini, F., Colaone, F., Roucairol, C., Laborie, M., Croughs, T., Grouard-Vogel, G., Lauwerys, B.R., 2016. Interferon α kinoid induces neutralizing anti-interferon α antibodies that decrease the

expression of interferon-induced and B cell activation associated transcripts: analysis of extended follow-up data from the interferon α kinoid phase I/II study. Rheumatology 55, 1901–1905. doi:10.1093/rheumatology/kew262

Edgar, R., 2002. Gene Expression Omnibus: NCBI gene expression and hybridization array data repository. Nucleic Acids Research 30, 207–210. doi:10.1093/nar/30.1.207

Fabregat, A., Jupe, S., Matthews, L., Sidiropoulos, K., Gillespie, M., Garapati, P., Haw, R., Jassal, B., Korninger, F., May, B., Milacic, M., Roca, C.D., Rothfels, K., Sevilla, C., Shamovsky, V., Shorser, S., Varusai, T., Viteri, G., Weiser, J., Wu, G., Stein, L., Hermjakob, H., D'Eustachio, P., 2017. The Reactome Pathway Knowledgebase. Nucleic Acids Research 46, D649–D655. doi:10.1093/nar/gkx1132

Grayson, P.C., Eddy, S., Taroni, J.N., Lightfoot, Y.L., Mariani, L., Parikh, H., Lindenmeyer, M.T., Ju, W., Greene, C.S., Godfrey, B., Cohen, C.D., Krischer, J., Kretzler, M., Merkel, P.A., 2018. Metabolic pathways and immunometabolism in rare kidney diseases. Annals of the Rheumatic Diseases 77, annrheumdis–2017–212935–1233. doi:10.1136/annrheumdis-2017-212935

Grayson, P.C., Steiling, K., Platt, M., Berman, J.S., Zhang, X., Xiao, J., Alekseyev, Y.O., Liu, G., Monach, P.A., Kaplan, M.J., Spira, A., Merkel, P.A., 2015. Defining the nasal transcriptome in granulomatosis with polyangiitis (Wegener's). Arthritis Rheumatol 67, 2233–2239. doi:10.1002/art.39185

Johnson, W.E., Li, C., Rabinovic, A., 2007. Adjusting batch effects in microarray expression data using empirical Bayes methods. Biostatistics 8, 118–127. doi:10.1093/biostatistics/kxj037

Kolesnikov, N., Hastings, E., Keays, M., Melnichuk, O., Tang, Y.A., Williams, E., Dylag, M., Kurbatova, N., Brandizi, M., Burdett, T., Megy, K., Pilicheva, E., Rustici, G., Tikhonov, A., Parkinson, H., Petryszak, R., Sarkans, U., Brazma, A., 2014. ArrayExpress update—simplifying data submissions. Nucleic Acids Research 43, D1113–D1116. doi:10.1093/nar/gku1057

Lauwerys, B.R., Hachulla, E., Spertini, F., Lazaro, E., Jorgensen, C., Mariette, X., Haelterman, E., Grouard-Vogel, G., Fanget, B., Dhellin, O., Vandepapelière, P., Houssiau, F.A., 2013. Down-regulation of interferon signature in systemic lupus erythematosus patients by active immunization with interferon α-kinoid. Arthritis Rheum 65, 447–456. doi:10.1002/art.37785

Malone, J., Holloway, E., Adamusiak, T., Kapushesky, M., Zheng, J., Kolesnikov, N., Zhukova, A., Brazma, A., Parkinson, H., 2010. Modeling sample variables with an Experimental Factor Ontology. Bioinformatics 26, 1112–1118. doi:10.1093/bioinformatics/btq099

Mao, W., Harmann, B., Sealfon, S.C., Zaslavsky, E., Chikina, M., 2017. Pathway-Level Information ExtractoR (PLIER) for gene expression data. bioRxiv 116061. doi:10.1101/116061

McKinney, E.F., Lee, J.C., Jayne, D.R.W., Lyons, P.A., Smith, K.G.C., 2015. T-cell exhaustion, co-stimulation and clinical outcome in autoimmunity and infection. Nature 523, 612–616. doi:10.1038/nature14468

Newman, A.M., Liu, C.L., Green, M.R., Gentles, A.J., Feng, W., Xu, Y., Hoang, C.D., Diehn, M., Alizadeh, A.A., 2015. Robust enumeration of cell subsets from tissue expression profiles. Nat Methods 12, 453–457. doi:10.1038/nmeth.3337

Obermoser, G., Presnell, S., Domico, K., Xu, H., Wang, Y., Anguiano, E., Thompson-Snipes, L., Ranganathan, R., Zeitner, B., Bjork, A., Anderson, D., Speake, C., Ruchaud, E., Skinner, J., Alsina, L., Sharma, M., Dutartre, H., Cepika, A., Israelsson, E., Nguyen, P., Nguyen, Q.-A., Harrod, A.C., Zurawski, S.M., Pascual, V., Ueno, H., Nepom, G.T., Quinn, C., Blankenship, D., Palucka, K., Banchereau, J., Chaussabel, D., 2013. Systems Scale Interactive Exploration Reveals Quantitative and Qualitative Differences in Response to Influenza and Pneumococcal Vaccines. Immunity 38, 831–844. doi:10.1016/j.immuni.2012.12.008

Oswald, M., Curran, M.E., Lamberth, S.L., Townsend, R.M., Hamilton, J.D., Chernoff, D.N., Carulli, J., Townsend, M.J., Weinblatt, M.E., Kern, M., Pond, C.M., Lee, A., Gregersen, P.K., 2015. Modular analysis of peripheral blood gene expression in rheumatoid arthritis captures reproducible gene expression changes in tumor necrosis factor responders. Arthritis Rheumatol 67, 344–351. doi:10.1002/art.38947

Pan, S.J., Yang, Q., 2010. A Survey on Transfer Learning. IEEE Transactions on Knowledge and Data Engineering 22, 1345–1359. doi:10.1109/TKDE.2009.191

Rovenský, J., Payer, J. (Eds.), 2009. Systemic lupus erythematosus disease activity index (SLEDAI), in: Dictionary of Rheumatology. Springer Vienna, Vienna, pp. 209–209. doi:10.1007/978-3-211-79280-3_1086

Sierra-Filardi, E., Nieto, C., Dominguez-Soto, A., Barroso, R., Sanchez-Mateos, P., Puig-Kroger, A., Lopez-Bravo, M., Joven, J., Ardavin, C., Rodriguez-Fernandez, J.L., Sanchez-Torres, C., Mellado, M., Corbi, A.L., 2014. CCL2 Shapes Macrophage Polarization by GM-CSF and M-CSF: Identification of CCL2/CCR2-Dependent Gene Expression Profile. The Journal of Immunology 192, 3858–3867. doi:10.4049/jimmunol.1302821

Stein-O'Brien, G.L., Arora, R., Culhane, A.C., Favorov, A., Garmire, L.X., Greene, C., Goff, L.A., Li, Y., Ngom, A., Ochs, M.F., Xu, Y., Fertig, E.J., 2017a. Enter the matrix: factorization uncovers knowledge from omics. doi:10.1101/196915

Stein-O'Brien, G.L., Carey, J.L., Lee, W.S., Considine, M., Favorov, A.V., Flam, E., Guo, T., Li, S., Marchionni, L., Sherman, T., Sivy, S., Gaykalova, D.A., McKay, R.D., Ochs, M.F., Colantuoni, C., Fertig, E.J., 2017b. PatternMarkers & GWCoGAPS for novel data-driven biomarkers via whole transcriptome NMF. Bioinformatics 33, 1892–1894. doi:10.1093/bioinformatics/btx058

Subramanian, A., Tamayo, P., Mootha, V.K., Mukherjee, S., Ebert, B.L., Gillette, M.A., Paulovich, A., Pomeroy, S.L., Golub, T.R., Lander, E.S., Mesirov, J.P., 2005. Gene set enrichment analysis: a knowledge-based approach for interpreting genome-wide expression profiles. Proceedings of the National Academy of Sciences 102, 15545–15550. doi:10.1073/pnas.0506580102

Tan, J., Doing, G., Lewis, K.A., Price, C.E., Chen, K.M., Cady, K.C., Perchuk, B., Laub, M.T., Hogan, D.A., Greene, C.S., 2017. Unsupervised Extraction of Stable Expression Signatures from Public Compendia with an Ensemble of Neural Networks. Cell Systems 5, 63–71.e6. doi:10.1016/j.cels.2017.06.003

Tan, J., Hammond, J.H., Hogan, D.A., Greene, C.S., 2016. ADAGE-Based Integration of Publicly Available Pseudomonas aeruginosa Gene Expression Data with Denoising Autoencoders Illuminates Microbe-Host Interactions. mSystems 1, e00025–15. doi:10.1128/msystems.00025-15

Thompson, J.A., Tan, J., Greene, C.S., 2016. Cross-platform normalization of microarray and RNA-seq data for machine learning applications. PeerJ 4, e1621. doi:10.7717/peerj.1621

Way, G.P., Greene, C.S., 2017. Extracting a Biologically Relevant Latent Space from Cancer Transcriptomes with Variational Autoencoders. bioRxiv 174474. doi:10.1101/174474

Welcher, A.A., Boedigheimer, M., Kivitz, A.J., Amoura, Z., Buyon, J., Rudinskaya, A., Latinis, K., Chiu, K., Oliner, K.S., Damore, M.A., Arnold, G.E., Sohn, W., Chirmule, N., Goyal, L., Banfield, C., Chung, J.B., 2015. Blockade of Interferon-γ Normalizes Interferon-Regulated Gene Expression and Serum CXCL10 Levels in Patients With Systemic Lupus Erythematosus. Arthritis & Rheumatology 67, 2713–2722. doi:10.1002/art.39248

Zelaya, R.A., Wong, A.K., Frase, A.T., Ritchie, M.D., Greene, C.S., 2016. Tribe: The collaborative platform for reproducible web-based analysis of gene sets. doi:10.1101/055913

Zhu, Y., Stephens, R.M., Meltzer, P.S., Davis, S.R., 2013. SRAdb: query and use public next-generation sequencing data from within R. BMC Bioinformatics 14, 19. doi:10.1186/1471-2105-14-19

Table 1. Datasets included in the systemic lupus erythematosus whole blood dataset.
Raw data availability refers to whether or not raw data was associated with the accession on ArrayExpress. We note any information that is relevant to the evaluations we performed.

Accession	Number of samples	Platform	Raw data availability	Publication
E-GEOD-65391*	996	Illumina HumanHT-12 v4	No	Banchereau et al., 2016
E-GEOD-11907	4	Affymetrix HGU133A and B	Yes	Chaussabel et al., 2008
E-GEOD-39088**	142	Affymetrix HGU133plus2	Yes	Lauwerys et al., 2013
E-GEOD-72747	30	Affymetrix HGU133plus2	Yes	Ducreux et al., 2016
E-GEOD-49454	177	Illumina HumanHT-12 v4	No	Chiche et al., 2014
E-GEOD-61635	129	Affymetrix HGU133plus2	Yes	N/A
E-GEOD-78193***	125	Agilent 4x44K	No	Welcher et al., 2015
Notes: *Longitudinal pediatric cohort; Accession includes neutrophil counts **IFN-alpha-kinoid trial ***AMG 811 trial				

FIGURE LEGENDS

Figure 1. Overview of dataset-specific PLIER and MultiPLIER. Boxes with solid colored fills represent inputs to the model. White boxes with colored outlines represent model output. **(A)** PLIER (Mao et al., 2017) automatically extracts latent variables (LVs), shown as the matrix B, and their loadings (Z). We can train PLIER model for each of three datasets from different tissues, which results in three dataset-specific latent spaces. **(B)** PLIER takes as input a prior information/knowledge matrix C and applies a constraint such that some of the loadings (Z) and therefore some of the LVs capture biological signal in the form of curated pathways or cell type-specific gene sets. **(C)** Ideally, an LV will map to a single gene set or a group of highly related gene sets to allow for easy interpretation of the model. PLIER applies a penalty on U to facilitate this. Purple fill in a cell indicates a non-zero value and a darker purple indicates a higher value. We show an undesirable U matrix in the top toy example (Ci) and a favorable U matrix in the bottom toy example (Cii). **(D)** If models have been trained on individual datasets, we may be required to find "matching" LVs in different dataset- or tissue-specific models using the loadings (Z) from each model. Using a metric like the Pearson correlation between loadings, we may or may not be able to find a well-correlated match between datasets. **(E)** The MultiPLIER approach: train a PLIER on a large collection of uniformly processed data from many different biological contexts and conditions (recount2; Collado-Torres et al., 2017)—a MultiPLIER model—and then project the individual datasets into the MultiPLIER latent space. The hatched fill indicates the sample dataset of origin. **(F)** Latent variables from the MultiPLIER model can be tested for differential expression between disease and controls in multiple tissues.

Figure 2. A PLIER model trained on a systemic lupus erythematosus (SLE) whole blood (WB) compendium learns SLE pathology-relevant LVs and divides biological signal and technical noise. (A) Selected LVs from the SLE WB PLIER U matrix. Purple fill in a cell indicates a non-zero value and a darker purple indicates a higher value. Only pathways with AUC > 0.75 in displayed LVs are shown. Panels B-D display the first two PCs from different subsets of the latent space or B matrix. Points are samples. Samples are colored by dataset of origin and datasets from the same platform manufacturer are similar colors. **(B)** PC1 and PC2 from PCA on the entire B matrix illustrates a platform- or dataset-specific effect. **(C)** PC1 and PC2 from only pathway- or geneset-associated LVs (LVs with AUC > 0.75 for at least one geneset) show a reduction in the technical variance evident in panel B. **(D)** PC1 and PC2 from only LVs that do not have an association with an input geneset (all AUC <= 0.75) show a similar pattern to that of all LVs. The dataset-specific effect in panels B and D is also observed in PCA on the gene-level gene expression data (Fig. S1).

Figure 3. MultiPLIER learns a neutrophil-associated LV that is well-correlated with neutrophil counts or estimates in multiple tissues and disease contexts. (A) SLE WB LV87, LV87 from the PLIER model trained on the entire SLE WB compendium, can predict neutrophil count in Banchereau, et al. dataset (Banchereau et al., 2016) despite being entirely unsupervised with respect to this task. **(B)** The MultiPLIER neutrophil-associated performs slightly better than the SLE WB-specific model. Here, the SLE WB compendium is projected into the multiPLIER latent space. Recall that the MultiPLIER model has not been exposed to the *specific* technical variance found in the SLE WB compendium. **(C)** MultiPLIER performance is not limited to whole blood, as it is highly correlated with neutrophil estimate from a state-of-the-art method for estimating immune infiltrate in solid and non-cancerous tissues, MCPcounter (Becht et al., 2016) in the NARES dataset (Grayson et al., 2015). NARES is a nasal brushing microarray dataset that

includes patients with ANCA-associated vasculitis, patients with sarcoidosis and healthy controls among other groups and was projected into the MultiPLIER latent space.

Figure 4. MultiPLIER learns type I and type II interferon latent variables that are perturbed in a manner consistent with interferon blockade. (A) Patients can be stratified into IFN-positive and IFN-negative groups using baseline expression of LV116. The expression of type II LV, LV140, does not show a clear IFN-negative and IFN-positive group. **(B)** Treated IFN-positive patients show a statistically significant reduction in LV116 expression at day 112 as compared to the placebo group (p-value from Wilcoxon rank sum test).

Figure 5. Subsampling of the recount2 compendium demonstrates the contribution of both sample size and breadth of biological conditions to PLIER model characteristics. The boxplot and points in black represent 10 repeats for a subsampling of the recount2 compendium (n =1640, same size as SLE WB compendium) and the whole SLE WB compendium. The blue diamonds are the values from the full recount2 PLIER model. **(A)** The number of latent variables (k) in a model. **(B)** The proportion of pathways supplied as input to the model that are significantly associated (FDR < 0.05) with at least one LV, termed pathway coverage, is similar between subsampled and SLE models and about half of the pathway coverage observed in the full recount2 PLIER model. This suggests that pathway coverage is constrained by sample size. **(C)** The proportion of latent variables that are significantly associated (FDR < 0.05) with at least one pathway or gene set is higher in the subsampled models as compared to the SLE WB models.

Figure 6. The latent space of the MultiPLIER model (no vasculitis data in the training set) shows agreement with the PLIER model trained on the NARES nasal brushing dataset (n = 79) particularly in the case of LVs of biological significance. The best match for a NARES model LV was determined by identifying the LV in the MultiPLIER model with the most similar loadings (Pearson correlation of Z). NARES data was projected into the MultiPLIER latent space (B) and the Pearson correlation between B matrices (NARES, MultiPLIER) was calculated. Density plot of the LV expression value correlation coefficients between best match LVs, shown in gray, demonstrates a rightward shift from the values for all pairs of LVs between models (e.g., including random pairs of LVs), shown in white. The blue points along the bottom of the graph represent best match correlation values for LVs from the NARES model that are significantly associated (FDR < 0.05) with at least one gene set. This suggests that the values between LVs that capture biological signal are particularly likely to be preserved in the transfer learning (MultiPLIER) case. Figure S6 shows the correlation values between best match LVs alongside the correlation between loadings.

Figure 7. A MultiPLIER-learned LV associated with macrophages is differentially expressed in three tissues from ANCA-associated vasculitis (AAV) and shows increased expression in severe or active disease. Differentially expressed LVs were identified by comparing all patient groups and using Benjamini-Hochberg correction (FDR). LVs with FDR < 0.05 were considered to be differentially expressed. MultiPLIER LV10 had an FDR < 0.05 in all three cohorts. **(A-C)** Jitter plots of MultiPLIER LV10 in three different tissues: nasal brushings (NARES dataset), kidney microdissected glomeruli, and peripheral blood mononuclear cells. Points and bars represent mean ± 2 * SEM. P-values are from a Wilcoxon rank sum test comparing the control group to the AAV group considered to have the most severe or active disease in the cohort. **(D)** The loadings of the top (highest weight) 25 genes for MultiPLIER LV10.

SUPPLEMENTAL FIGURE LEGENDS

Figure S1. Principal components 1-5 for the SLE WB expression compendium. Related to Figure 2. Points are samples. Samples are colored by dataset of origin and datasets from the same platform manufacturer are similar colors. PC1-5 account for 0.502 of the variance explained.

Figure S2. MultiPLIER learns a plasma cell signature that is elevated in more severe disease. Related to Figure 3. (A) Selected LVs from the MultiPLIER U matrix. Purple fill in a cell indicates a non-zero value and a darker purple indicates a higher value. Only pathways with FDR < 0.05 in displayed LVs are shown. We selected LVs discussed in the main text. **(B-C)** In Banchereau, et al. citation the authors stratified patients into three disease activity groups—DA1 (SLEDAI: 0-2), DA2 (SLEDAI: 3-7), and DA3 (SLEDAI > 7)—and demonstrated that patients with more severe disease had higher plasmablast counts (via FACS). We did not provide a plasmablast-specific gene set to the PLIER model during training. However, plasma cell from signatures adapted from the Immune Response In Silico (IRIS) project (Abbas et al., 2005) and Newman, et al. (Newman et al., 2015)(included in the PLIER R package). A plasmablast is the precursor to a plasma cell and we therefore assume that these cell types share gene expression patterns. Both the SLE WB **(B)** and MultiPLIER **(C)** models learn an LV associated with this signature (SLE WB LV52 and MultiPLIER LV951, respectively). P-values are from a pairwise t-test with Bonferroni correction. Both LVs show a similar pattern to the original publication FACS plasmablast counts. However, there is evidence that the MultiPLIER LV may more specific to plasma cells, as SLE WB LV52 is also significantly associated (FDR < 0.05) with REACTOME Unfolded Protein Response and REACTOME Asparagine N-linked Glycosylation (Fabregat et al., 2017)(Figure 2A). MultiPLIER LV951 is only significantly associated with plasma cell gene sets (data not shown).

Figure S3. Projection of isolated leukocyte gene expression data, as measured by microarray, into MultiPLIER latent space demonstrates that the model captures additional cell type-specific signatures. Related to Figure 3. Heatmap of a subset of the B matrix for E-MTAB-2452 (McKinney et al., 2015), a dataset comprised of isolated immune subsets (CD4+ T cells, CD14+ monocytes, and CD16+ neutrophils) from patients with autoimmune diseases. Latent variables (rows) were selected if their names contained one of the following terms: CD4, monocyte, or neutrophil. The sample (column) annotation bar is colored based on the cell subset of that sample. The LV (row) annotation bar is colored by the term associated with the LV.

Figure S4. IFN-K expression data in the SLE WB PLIER model and modular transcriptional repertoire framework. Related to Figure 4. Chiche, et al. (Chiche et al., 2014) identified three IFN-related modules—M1.2, M3.4, and M5.12—and demonstrated that M1.2 was strongly induced by type I IFN, whereas the other two modules were similarly responsive to type I and type II IFN. Type I module M1.2, and to a lesser extent M3.4, from the whole blood modular framework divides patients into IFN-positive and IFN-negative groups. Expression for each module was summarized using the mean expression value of all genes in the module. **(B)** Treated IFN-positive patients show a statistically significant reduction in M1.2 expression at day 112 as compared to the placebo group (p-value from Wilcoxon rank sum test). **(C)** LV6, LV69, and LV110 from the SLE WB PLIER model were all significantly associated with type I IFN (REACTOME Interferon alpha/beta signaling). Patients were classified as IFN-negative if they had one of the nine lowest expression values in 2 out of 3 of these LVs (typically LV6 and LV110). Recall that these samples are in the training data for this model. **(D)** Treated IFN-positive patients do not show a statistically significant decrease in LV6 expression at day 112 as compared to placebo (p-

value from Wilcoxon rank sum test). SLE WB LV6 was selected for display and comparison as it had the highest AUC for REACTOME Interferon alpha/beta signaling in the SLE WB model.

Figure S5. Expression patterns from trial of type II IFN blockade agent AMG 811 in the MultiPLIER latent space, whole blood modular framework, and SLE WB PLIER. Related to Figure 4. The type II IFN LV from the MultiPLIER model, LV140, shows a trend towards decreased expression during days 15 and 56 of the AMG811 trial (top panel). The modules most likely to capture type II IFN expression, M3.4 and M5.12, show a similar trend (middle pattern). The LVs or modules indicative of type I signaling show no such pattern (MultiPLIER LV116, whole blood modular framework M1.2, and all SLE WB LVs). No SLE WB PLIER LV was significantly associated with type II (as determined by the lack of association with the REACTOME Interferon gamma signaling pathway).

Figure S6. Comparison of the loadings (Z) and latent space (B) of the NARES PLIER model to the MultiPLIER loadings and the projection of the NARES dataset into the MultiPLIER latent space. Related to Figure 6. Heatmap of correlation values between the loadings and latent space of best match LVs, where rows are the LVs from the NARES PLIER model. A more saturated red color indicates a higher correlation value. A NARES LV's best match is determined by calculating the Pearson correlation between loadings (Z matrices) and selecting the LV from the MultiPLIER model with the highest correlation to that NARES LV. The columns are correlation values for the loadings (Z) and the LV expression (B) between that NARES LV (row) and its best match in the MultiPLIER model. The row annotation bar indicates whether or not a NARES LV is significantly associated (FDR < 0.05) with one or more gene sets supplied during training.

Figure S7. MultiPLIER learns a latent variable that captures the activity of antigens in antineutrophil cytoplasmic antibody (ANCA)-associated vasculitis despite no vasculitis training data. Related to Figure 7. The peripheral blood mononuclear cell (PBMC) fraction data from Cheadle, et al. (Cheadle et al., 2010) was projected into the MultiPLIER latent space. The authors of Cheadle, et al. identified three groups: GPA-positive ("WG-positive" for Wegener's Granulomatosis in the original 2010 publication), GPA-negative ("WG-negative" in the original publication), and controls. **(A)** MultiPLIER LV599 is differentially expressed between the three groups (FDR = 1.17e-08). Points and bars represent mean ± 2 * SEM. **(B)** The loadings of the top (highest weight) 50 genes for MultiPLIER LV599 includes ANCA antigens, shown with black bars. The original publication reported differential expression of ANCA antigens in the PBMC fraction.

Figure S8. A MultiPLIER-learned LV associated with the HIF-1a transcription factor network is differentially expressed in three tissues from ANCA-associated vasculitis (AAV) and shows increased expression in severe or active disease. Related to Figure 7. Differentially expressed LVs were identified by comparing all patient groups and using Benjamini-Hochberg correction (FDR). LVs with FDR < 0.05 were considered to be differentially expressed. MultiPLIER LV937 had an FDR < 0.05 in all three cohorts. **(A)** Venn Diagram of LVs differentially expressed in three tissues. **(B-D)** Jitter plots of MultiPLIER LV937 in three different tissues: nasal brushings (NARES dataset), kidney microdissected glomeruli, and peripheral blood mononuclear cells. Points and bars represent mean ± 2 * SEM. P-values are from a Wilcoxon rank sum test comparing the control group to the AAV group considered to have the most severe or active disease in the cohort. **(E)** The loadings of the top (highest weight) 25 genes for MultiPLIER LV937.

Figure 1

Figure 2

Figure 3

MultiPLIER

A Number of LVs

B Pathway Coverage

C LVs significantly associated with pathways

Figure 5

Figure 6

Figure S1

Figure S2

Color Key

−0.4   0   0.4

Value

Isolated cell populations microarray

CD4 T cell
CD14, monocyte
CD16, neutrophil

603,SVM Neutrophils
985,IRIS_Neutrophil−Resting
7,IRIS_Neutrophil−Resting
975,SVM T cells CD4 memory resting
2,SVM Monocytes
942,SVM T cells CD4 memory activated
823,SVM T cells CD4 memory resting
57,SVM T cells CD4 memory activated
675,IRIS_Neutrophil−Resting
953,IRIS_Monocyte−Day1
649,IRIS_Monocyte−Day1
524,IRIS_Neutrophil−Resting
854,IRIS_Monocyte−Day1

recount2 LVs

Samples

Figure S3

Whole Blood Modular Framework

A Baseline

B M1.2 Day 112 of IFN−K treatment

SLE WB PLIER

C Baseline

D LV6 Day 112 of IFN−K treatment

significant association with one or more pathways (FDR < 0.05)

0.12	0.33	NARES LV1
0.15	0.06	NARES LV2
0.71	0.96	NARES LV3
0.14	0.93	NARES LV4
0.15	0.34	NARES LV5
0.17	0.7	NARES LV6
0.19	0.43	NARES LV7
0.51	0.77	NARES LV8
0.1	0.27	NARES LV9
0.22	0.6	NARES LV10
0.21	0.7	NARES LV11
0.2	0.6	NARES LV12
0.21	0.58	NARES LV13
0.5	0.93	NARES LV14
0.38	0.69	NARES LV15
0.2	0.74	NARES LV16
0.34	0.9	NARES LV17
0.17	0.37	NARES LV18
0.11	0.46	NARES LV19
0.09	0.3	NARES LV20
0.35	0.5	NARES LV21
0.55	0.51	NARES LV22
0.07	0.15	NARES LV23
0.07	0.11	NARES LV24
0.41	0.67	NARES LV25
0.16	0.68	NARES LV26
0.39	0.68	NARES LV27
0.37	0.87	NARES LV28
0.24	0.78	NARES LV29
0.42	0.73	NARES LV30
0.44	0.87	NARES LV31
0.1	0.35	NARES LV32
0.26	0.61	NARES LV33
0.09	0.16	NARES LV34

Figure S6

Figure S8