



OPEN

Interpretable machine learning identifies paediatric Systemic Lupus Erythematosus subtypes based on gene expression data

Sara A. Yones¹ , Alva Annett¹, Patricia Stoll², Klev Diamanti³ , Linda Holmfeldt³ , Carl Fredrik Barrenäs¹, Jennifer R. S. Meadows^{4,8}  & Jan Komorowski^{1,5,6,7,8} 

Transcriptomic analyses are commonly used to identify differentially expressed genes between patients and controls, or within individuals across disease courses. These methods, whilst effective, cannot encompass the combinatorial effects of genes driving disease. We applied rule-based machine learning (RBML) models and rule networks (RN) to an existing paediatric Systemic Lupus Erythematosus (SLE) blood expression dataset, with the goal of developing gene networks to separate low and high disease activity (DA1 and DA3). The resultant model had an 81% accuracy to distinguish between DA1 and DA3, with unsupervised hierarchical clustering revealing additional subgroups indicative of the immune axis involved or state of disease flare. These subgroups correlated with clinical variables, suggesting that the gene sets identified may further the understanding of gene networks that act in concert to drive disease progression. This included roles for genes (i) induced by interferons (*IFI35* and *OTOF*), (ii) key to SLE cell types (*KLRB1* encoding CD161), or (iii) with roles in autophagy and NF- κ B pathway responses (*CKAP4*). As demonstrated here, RBML approaches have the potential to reveal novel gene patterns from within a heterogeneous disease, facilitating patient clinical and therapeutic stratification.

Paediatric systemic lupus erythematosus (pSLE) is a rare, clinically and genetically heterogeneous systemic autoimmune disease with a prevalence of between 3.3 and 8.8 per 100,000 children¹. The disease course is unpredictable, with periods of remission and flares that lead to cumulative damage over time². SLE is classified by the presence of at least 4 out of 11 of clinical criteria³, with disease activity (DA) severity calculated based on composite scores, including Systemic Lupus Erythematosus Disease Activity Index (SLEDAI)⁴. Genome-wide association studies have identified more than 130 SLE-associated loci⁵, including those driven by interferons⁶, or those controlling inflammation and tissue response to injury⁷. Together these have been used to highlight the link between SLE and viral responses⁸. However, the trigger that initiates the expression of these genes and the progression of SLE disease remains poorly understood⁹.

Efforts to unravel the SLE gene expression pathway have been initiated. A 2016 study of paediatric disease examined the personal transcriptomic profiles of 158 patients using linear mixed models built on blood expression data from 15,386 transcripts¹⁰. The transcript panel utilised for this process considered each gene locus individually, and correlated the binary up- or down-regulation patterns with patient phenotypes. The result was the stratification of patients into distinct subclasses, with an enrichment of neutrophil expressed transcripts noted as a patient passed from the low DA1 state to the high DA3 form of disease. While the molecular pathways proposed by the study have led to a better understanding of personal disease progression¹⁰, the analysis lacked the co-predictive power of rule-based machine learning (RBML) models.

¹Science for Life Laboratory, Department of Cell and Molecular Biology, Uppsala University, Uppsala, Sweden. ²Department of Biosystems Science and Engineering, ETH Zurich, Zurich, Switzerland. ³Science for Life Laboratory, Department of Immunology, Genetics and Pathology, Uppsala University, Uppsala, Sweden. ⁴Science for Life Laboratory, Department of Medical Biochemistry and Microbiology, Uppsala University, Uppsala, Sweden. ⁵Washington National Primate Research Center, Seattle, USA. ⁶Swedish Collegium for Advanced Study, Uppsala, Sweden. ⁷The Institute of Computer Science, Polish Academy of Sciences, Warsaw, Poland. ⁸These authors contributed equally: Jennifer R. S. Meadows and Jan Komorowski. ✉email: sara.younes@icm.uu.se; jan.komorowski@icm.uu.se

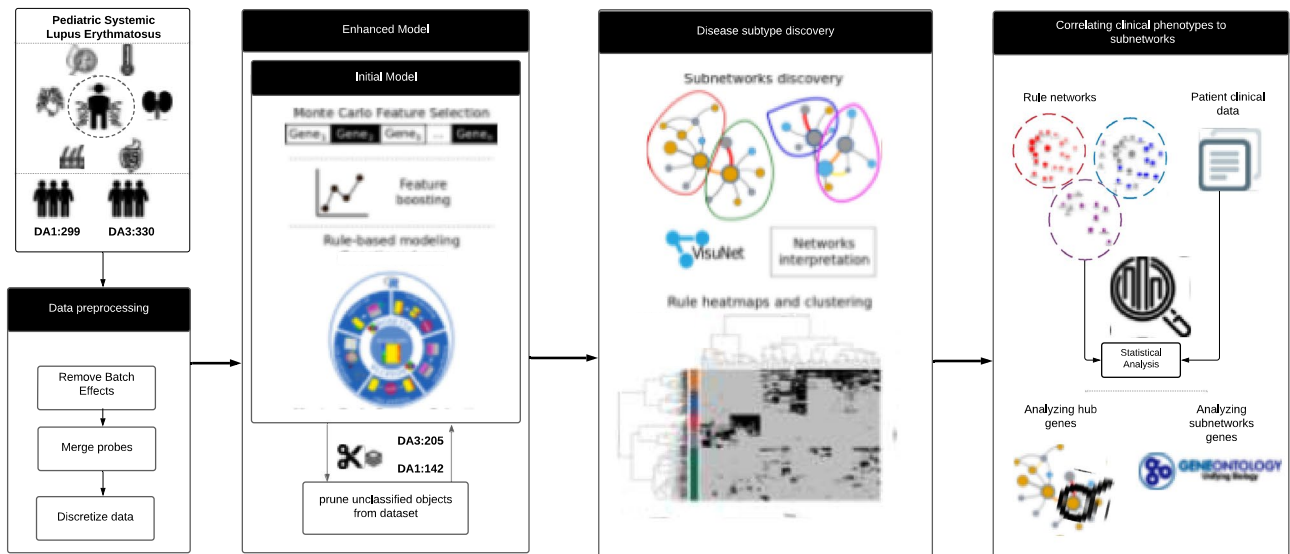


Figure 1. Overview of the modelling process implemented to classify and interrogate gene expression relationships between DA1 and DA3.

Machine learning (ML) approaches are well suited to address this process, as they can model and characterise data with very high dimensionality, such as that generated through personal transcriptomics. However, the majority of methods work as black boxes. These offer little to no explanation in terms of how, and why, a specific classification decision is made. For clinical -omics, understanding how a classification decision is made, may offer insight into the underlying biological mechanisms, for example contrasting a disease state to healthy controls¹⁰. Interpretable ML methods such as RBML models, offer classification transparency^{11,12}. We applied RBML that is based on rough set theory¹³. It uses Boolean reasoning to identify the minimal set of features that can discern decision classes (reducts). Reducts are subsequently overlaid onto transcriptomics data samples to create IF-THEN rules. One of the main advantages of this method is co-prediction, i.e., the identification of descriptors that collaboratively correctly classify samples from the data into the outcomes. Co-prediction can provide insight into the candidate biological processes beyond of what can be learnt through co-expression networks.

In the current study, we apply a RBML approach using rough sets to existing pSLE blood transcriptome data¹⁰. Here, the goal was to identify the genes and interactions that demarcate a low pSLE DA1 state from a high DA3 state. The disease sub-groups discovered were intersected with available clinical data, revealing gene sets key to the progression of disease and the involvement of the innate and acquired immune arms. These genes, and their protein products, have the potential to be translated to biomarkers, or could be suggested points for therapeutic intervention.

Results

Minimum gene set model discerns DA1 from DA3. The two extremes of disease activity were defined with SLEDAI score (DA1: SLEDAI=0–2; DA3: SLEDAI>7). The initial rule-based model was built with R.ROSETTA¹⁴ using data from 629 unique patient clinical visits (observations) and the discretised gene expression value for each DA1 and DA3 patient visit (features: 33,006 probes for 629 observations; Fig. 1). This initial model had an overall prediction accuracy of 71% using tenfold cross validation (Supplementary Fig. S1 online). The observations (visits) incorrectly classified by the model (Supplementary Fig. S2 online) were pruned to achieve a better separation between DA1 and DA3 then intersected with the patient metadata in order to understand the potential reasons behind their misclassification. Observations were more likely to be pruned or removed based on patient treatment, low SLEDAI score or the number of days since diagnosis (Logistic regression p-value for all <0.05; Supplementary Fig. S3 online; Supplementary Table S1). No significant association was observed between removed observations and clinical symptoms, and the significant associations were a reflection of observations removed from class DA3 (38%, 125/330 removed) rather than reductions from DA1 (53%, 157/299 removed).

Following Monte Carlo Feature selection (MCFS)¹⁵ on the pruned dataset, 4980 genes were available and subsequently used to build an enhanced rule-based model. Gene set enrichment analysis revealed terms connected to neutrophils (e.g., activation, mediation, degranulation) and the production and degradation of gene products (e.g., transcription initiation and nonsense-mediated decay; Supplementary Fig. S4 online). This suggests a difference in neutrophil mediated immune response between patients with DA1 and DA3, a known functional shift in SLE manifestation between disease states⁴.

Feature boosting was performed to identify the optimal number of genes for the model (Fig. 1). Empirical studies revealed that model accuracy was lost if more than 200 of the top 4980 MCFS ranked genes were used for this process (Supplementary Fig. S5 online). Iterative R. ROSETTA computational rounds added genes from the starting set of 200, with maximum model accuracy of 81% achieved with a minimum set of 34 genes (Fig. 2;

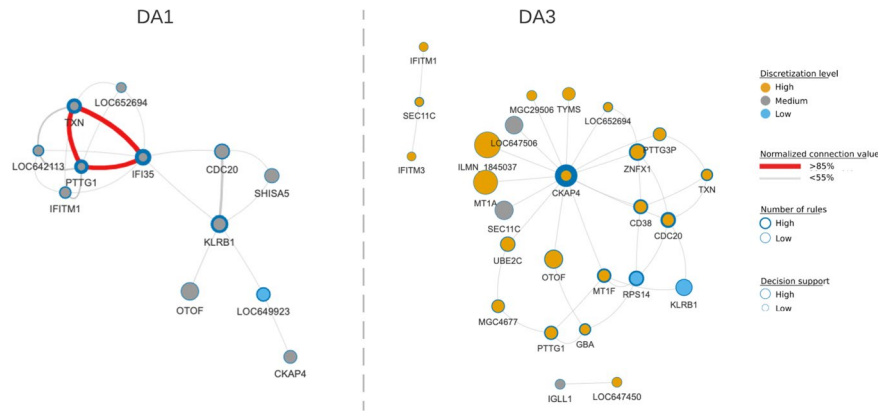


Figure 2. The rule networks discern the disease states. DA1 is largely defined by medium gene expression, whereas DA3 includes more genes, and those that were highly expressed. For each decision class, internal node colour indicates discretised gene expression value (high, medium, low; orange, grey, blue), node size is proportional to the number of objects supporting rules associated to a node, node border thickness is proportional to the number of rules associated to a node (low, high; circle border thin, thick) and edges connecting nodes represent normalised connection values ($< 55\%$, $\geq 85\%$; grey, red with increasing line thickness per support interval). The latter is the strength of the co-appearance of connected nodes in rules supporting a decision class. The network was filtered to visualise rules with minimum support of 10% and rule p -value ≤ 0.05 .

Supplementary Table S2 online). These genes were used in 22 and 44 classifying rules for DA1 and DA3 respectively. The model mirrored the structure of the initial model (Supplementary Fig. S1 online). Figure 2 shows DA1 and DA3 were again split, however with a reduction of complexity, in terms of rules (edges) connecting the genes (nodes) and a refinement of the central hub genes. The 10% points gain in the model accuracy provided improvement in terms of a clearer and visible separation between the disease activities in the rule networks (RN); this gain in accuracy was too small to imply an overfitting of the model. The similarity between the network of the initial model and the enhanced model implied that removed objects were unnecessary for classification of DA1 and DA3 since their removal did not significantly impact the main network structure or the rule model.

In DA3, hub gene *CKAP4* was surrounded by a thick blue border, indicating the importance of this gene to predicting this disease state. In fact, *CKAP4* was a member of 14/44 co-prediction rules (Supplementary Table S2 online). The protein product of this gene, CKAP4, formerly CLIMP63, can act to regulate endoplasmic reticulum (ER) nanodomain homeostasis via shaping the luminal space or through interaction with other ER-resident proteins¹⁶. *CKAP4* was highly expressed (orange), whereas connected gene *SEC11C* showed a medium level of discretised expression (grey), and *RPS14* was lowly expressed (blue). In DA1, *IFI35* and *KLRB1* were both hub genes with medium expression levels. However, the latter had larger number of observations supporting its membership to rules (larger node size) but contributed to slightly fewer rules than *IFI35* (thinner circle border size: *IFI35*, 6/22 rules; *KLRB1*, 4/22 rules). CD161/NKR-P1A, encoded by *KLRB1*, is a surface receptor of natural killer (NK) cells and subtypes of T lymphocytes, whereas *IFI35* encodes the Interferon-induced 35 kDa protein, a proinflammatory damage-associated molecular pattern (DAMP) molecule in the innate immune pathway¹⁷.

The membership of genes to the rule networks was not exclusive. For example, both *IFITM1*, which encodes interferon-induced transmembrane protein 1, and *KLRB1* appeared in DA1 and DA3 although with different expression values (Supplementary Table S2 online). The sharing of genes across rules was more common in DA1 (4/12 plotted genes are unique to DA1) than DA3, where 18/26 were unique to that class. The model showed that the type 1 interferon response term was limited to DA1 (*IFI35* and *IFITM1*) whereas B-cell activation was restricted to DA3 (*CD38* and *IGLL1*) (Supplementary Fig. S6 online). However, while each term was enriched based on very few genes, it should be noted that these genes were present in multiple rules.

Patient subgroups reflect clinical manifestations. Hierarchical clustering of the enhanced model results (i.e., the membership of observations for each rule) revealed five subgroups largely contained within DA1 (C1 and C2) or DA3 (C3, C4 and C5) (Fig. 3A). The model and sub-groups were tested for significance to confirm that they cannot be attained by using random data. The significance was tested using permutation of DA state (p -value ≤ 0.05). These sub-clusters were subsequently projected onto the RNs (Fig. 3B). Of note, the C1 and C2 sub-clusters were not restricted to the DA1 rule set, however C4 and C5 reflected partially intersecting networks that were all included in C3 and limited to DA3 (Fig. 3B). In comparison to C3, the DA3 hub gene *CKAP4* was absent from C4, whereas the two small unconnected DA3 networks were absent from C5. Due to the small number of genes available for consideration, a sub-cluster-based gene enrichment analysis was not informative for all sets. The C4 and C5 enrichments were largely based on the combination of two genes (*MT1F*, *MT1A*; 11 genes available) and suggested response to ion levels (see Supplementary Fig. S7 online), whereas C3 was again led by a small number of gene combinations (e.g., cell cycling and division: *CDC20*, *PTTG1*, *PTTG3P*, *UBE2C*; B-cell pathways: *CD38*, *GBA*, *TYM*) but this cluster also included the *MT1F*, and *MT1A* signals.

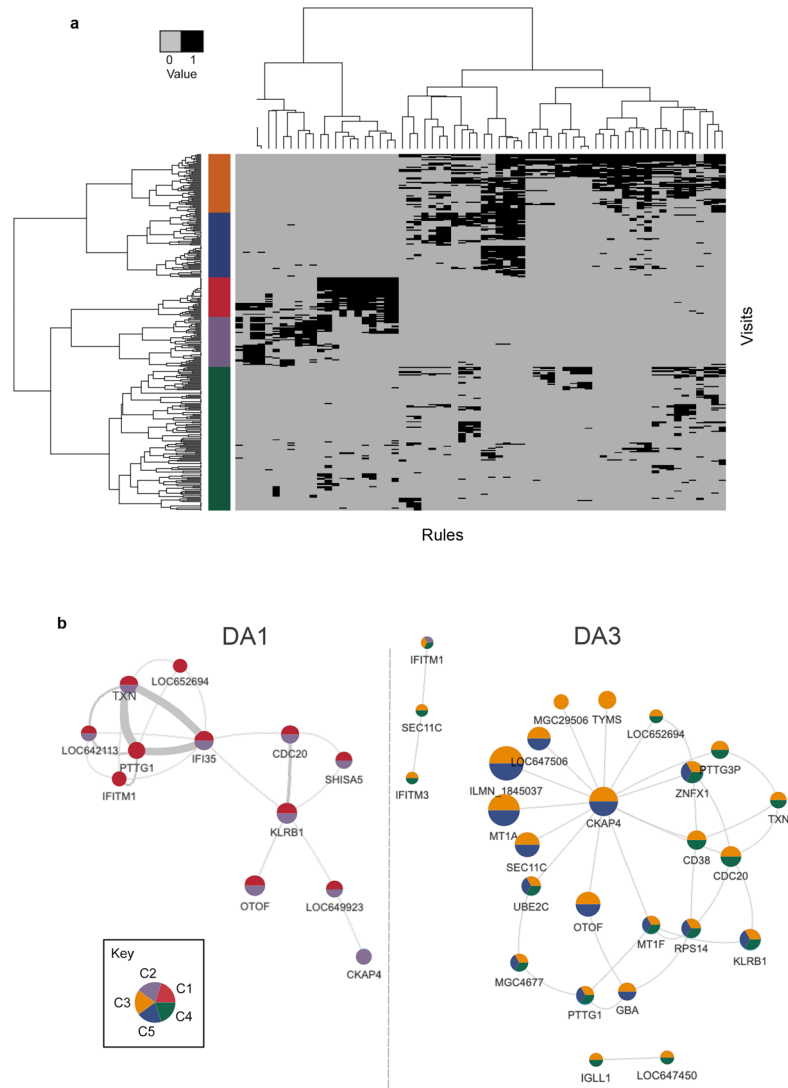


Figure 3. Hierarchical clustering of the model rules showed the major subdivision between the DA clusters. (a) Supported rules (black) and unsupported rules (grey) distinguish five disease subgroups that were projected into the (b) RN where group (cluster) membership is indicated by pie colour.

The relationship between clinical phenotype (Supplementary Table S3) and sub-cluster was explored in two ways. First, to assess clinical association to a sub-cluster, the phenotype values supporting that sub-cluster were compared to those that did not. To interrogate which rule(s) were driving that pattern, a similar assessment was performed, this time for visits supporting a rule within the sub-cluster. The examination of continuous phenotypes showed that these measures were only significantly different between the three DA3 clusters and not between the two DA1 clusters (Tukey HSD adjusted p -value < 0.05 ; Supplementary Table S4 online). However, for DA1, the C1 and C2 clusters did contain the majority of low SLEDAI score visits (~ 1.7 in each, Supplementary Table S4, Supplementary Fig. S8 online), with C1 tending towards lower alanine aminotransferase (ALT) and serum creatinine (CR) values compared with the C2 cluster. As expected, the DA3 cluster contained the higher SLEDAI scores (C4 ~ 8.8 , C5 ~ 12.1 , C3 ~ 14.6). C4 was largely reflective of low measures for anti-dsDNA antibody, erythrocyte sedimentation (ESR) and white blood cell count (WBC). C5 presented lower ALT and aspartate aminotransferase levels (AST), while C3 was most representative of active disease, with low complement factor C3 and C4 values (Supplementary Table S4, Supplementary Fig. S8 online). Only two phenotypes, lymphocyte percent (LP) and neutrophil percent (NP), were significantly different in all pairwise DA3 cluster comparisons. LP was highest in C4 and NP, highest in C5. C3 was intermediate for both (Supplementary Table S4, Supplementary Fig. S8 online).

In terms of categorical phenotypes, no significant association was detected between sex or race for each of the five clusters. In C1, the alopecia category was enriched when compared with all others (Fisher exact test p -value = 0.04; Supplementary Fig. S8 online). In C2, the musculoskeletal term and both oral steroid and nephritis treatment groups were enriched (all Fisher exact test p -value < 0.05). Treatment could not be ruled out as the factor driving differences between this and other clusters (Supplementary Fig. S9 online).

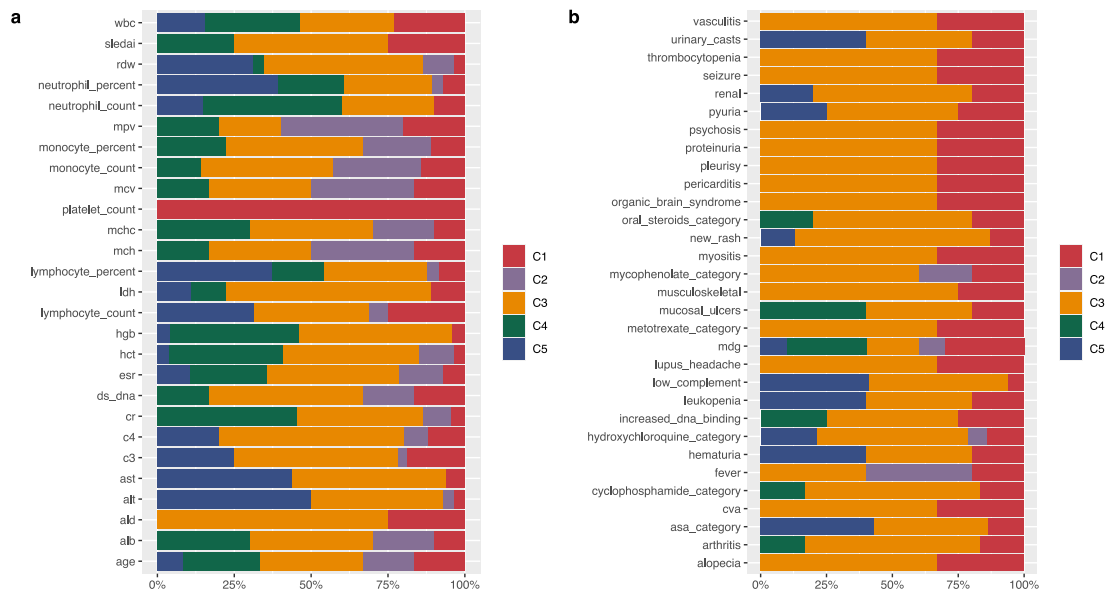


Figure 4. Fraction of rules per cluster significantly associated with (a) continuous and (b) categorical phenotypes. See Supplementary Table S3 online, for a list of clinical variables and phenotypes abbreviations.

Rules reveal which gene co-predictions drive phenotype correlation. To interrogate which genes and rules drove the phenotypic associations, a closer examination of the rules within the clusters was performed. To associate rules to the discovered clusters a frequency distribution was built for all rules with support set matching at least 10% of the visits assigned to each of the discovered clusters. Based on the distribution 20% match was an empirical threshold for assigning rules to each (Supplementary Fig. 10 online). Figure 4 illustrates the fraction of rules from each cluster that were significantly associated with a phenotype, either continuous or categorical. Overall, rules from C1 or C3 were significantly associated with all phenotypes displayed (Fig. 4), an enrichment not seen with the other clusters. Interestingly, whilst no individual continuous phenotype was significantly different between the two DA1 clusters, or categorical phenotype different between the DA3 clusters, the graphs clearly showed that the same was not true for the proportion of rules significantly associated with a phenotype in either class.

For example, in the continuous class, rules from both DA1 clusters were significantly associated with lymphocyte count (LC; Fig. 4a). There, three unique rules were contributed by C1 (rules 5, 44, 56), whereas the fourth rule was shared by both clusters (rule 41: *KLRB1*, *SEC11C*; Supplementary Table S5 online). Interestingly from the seven genes contained across the four rules, only the gene encoding the signal peptidase complex catalytic subunit, *SEC11C*, showed decreased expression, all others had medium values. This maintenance of gene expression likely explained the overall lack of significant difference between clusters for this trait.

For the DA3 clusters, a significant difference was recorded for the complement factor C3 phenotype between the C3 cluster (mean 62.1 mg/dL) and the C5 cluster (mean 85.9 mg/dL) (Wilcoxon test p -value = 2.4×10^{-3} ; Supplementary Table S4 online). An examination of the rules associated with phenotype C3 revealed that 17 rules were significantly linked to this phenotype in cluster C3, whilst only eight were found in the C5 cluster (Supplementary Table S5 online). All C5 rules were shared with C3, and no rules were contributed from C4 (Fig. 4a). As expected from the associated RN, none of the nine rules unique to cluster C3 showed discrete gene membership, rather they served to illustrate how in comparison to C5, rules represented by network edges could introduce additional unique features that may serve to explain the phenotypic difference. For example, shared rule 15 (*CKAP4*, *MT1F*) can form an extended connection with rules 4 (*MT1F*, *KLRB1*), 23 (*CKAP4*, *SEC11C*) and 51 (*MT1F*, *PTTG1*), widening this network to include genes *KLRB1*, *SEC11C* and *PTTG1* (Supplementary Table S6 online). Each of these genes had previously been associated with SLE, but the link was not always clear. As noted before, *KLRB1*, expressed by NK cells and shown to be in the medium discretised expression level here, has been implicated in the regulation of the interferon gamma immune response¹⁸. *SEC11C*, encodes a subunit of microsomal signal peptidase complex and was the only DA3 gene maintained within medium levels for this phenotype. This gene was previously shown to be significantly down regulated in the T cells of adult SLE patients with low complement levels¹⁷. *PTTG1* was previously linked to SLE via SNP association¹⁹, although it was later shown that the risk allele was tagging the nearby microRNA, miR-146a, and this was down-regulated in European disease²⁰.

Discussion

The use of machine learning in the current study has served to identify the key regulatory networks that underlie two disease states, DA1 and DA3, of the highly heterogeneous condition, paediatric systemic lupus erythematosus (pSLE). In doing so, the high dimensionality of data drawn from 33,006 gene expression measures across 629 paediatric patient visits has been reduced to co-predictive networks linked via genes. These genes were under-represented or down-weighted in published studies of SLE differential gene expression (DGE) profiling

(Supplementary Fig. S11 online). The result here was five sub-networks; two distinguishing DA1, perhaps as a result of treatment response, and three subgroups not related to treatment, within the more severe DA3 disease state.

Two major factors underpinned the difference in the results observed here, versus those generated by others in the field. The first was the study of patient visits, rather than individuals over time via longitudinal gene expression. The second was methodological, as RNs are co-predictive and as such, are conceptually different from co-expression networks. The goal here was to delve into the co-predictive RNs based on gene expression at different stages of disease, potentially creating a set of biomarkers, which could be used to stratify patient subgroups for clinical trials or personalised medicine based on their disease state at a particular time. This contrasts to the prognostic goals of others using the same dataset^{10,21}.

Let us set the scene. For the transcriptomic data analysed here, the nodes of an RN are genes and their discretised expression values. The edges between two nodes of an RN are formed from pairs of genes and their discretised expression values as they co-occurred in the IF-part of rules (Fig. 2). Significantly, in one outcome a gene may have one discretised value, but in the other outcome it will have a different value. It follows that each outcome has its own network. As such, co-prediction can provide insight into the candidate biological processes characteristic of the given outcome. For example, one combination of descriptors, i.e., pairs of gene and their discretised value may be associated to DA1 state, and another pair to DA3. This is in contrast to co-expression networks that identify genes that are co-expressed, not necessarily co-predictive of the outcomes.

SLE is a condition that spans the axes of both autoinflammatory and autoimmune disease. In this study, three DA3 subgroups were identified. The C3 sub-group sits on the autoimmune side, and had the clinical hallmarks of hypocomplementemia (low C3 and C4 clinical measures) in combination with high anti-dsDNA values, whilst the C4 sub-group likely represented the autoinflammatory side, with normal complement levels and low anti-dsDNA values (Supplementary Table S4 online). This was reinforced by the higher SLEDAI scores observed in C3 versus C4. Cluster C5 likely represented the intermediate stage between C3 and C4, where a significant shift between neutrophil and lymphocyte involvement is observed. This could translate to an immune complex driven disease state in C5, where the type I interferon process was active (low lymphocyte percent and increased neutrophil involvement). In studies using independent patient groups, both changes in complement ratio (C3/C4)²² and the categorisation of neutrophil to lymphocyte ratio (NLR)²³, have been suggested as ways to distinguish SLE patient groups. Here, network analysis and unsupervised clustering combined both C3/C4 and NLR biomarker sets and resulted in three separate groups spanning these factors. The novelty in the current study lies in linking the clusters to co-predictive RNs, and this was the second major factor differentiating this work from others.

While the application of machine learning approaches to the big data sets generated by biology -omics is not new²⁴, the approach used here removes the 'black box' interpretation of both the modelling and the results. This is required in the trade-off between predictability and interpretability²⁵. Here we accepted the potentially reduced, but still high prediction accuracy of 81%, in favour of transparent classical models that perform well when the number of features available in the dataset (i.e. observations versus genes) outnumber the observations²⁶. It is important to note that the rough sets approach to constructing rules, is based on finding all minimal subsets of features that preserve discernibility of the decision classes from the original set. The rules will contain conjunctions of genes that may reflect different levels of gene regulation but that do not need to be co-expressed. In RNs, the genes and their regulation levels are associated to the outcome and discern the decision classes (here DA1 or DA3) based on the training data, while in co-expression networks the genes are co-expressed with other genes and may not discern the outcomes. The R.ROSETTA method used for constructing the model has been shown to outperform other existing rule based methods¹⁴, and has the key distinction of being the only method that can compute a significance level for the rules in the model. This is useful for calculating model prediction reliability, but it is the use of a minimum set of significant rules that served to highlight the genes contributing most strongly to the separate networks.

In practice, this was illustrated by the hub genes for DA1 (e.g., *IFI35*, *KLRB1*) and DA3 (e.g., *CKAP4*, *OTOF*; Fig. 2). *IFI35* expression is stimulated in response to IFN- α ²⁷ and it can act intracellularly as a negative switch in the innate immune pathway via retinoic acid-inducible gene I regulation²⁸. Extracellularly, the opposite effect has been observed, and the *IFI35* molecule can act as a DAMP, and serve to activate the NF- κ B pathway in macrophages via TLR4 signalling¹⁷. The end result is the release of proinflammatory cytokines, including interleukin 6 and tumour necrosis factor¹⁷. In DA1, *IFI35* expression is observed within the medium range, but a change in this value could be key in driving DA1 patients back to a remissive or inactive SLE state. Likewise, the maintained medium expression of *KLRB1* (encoding the surface receptor CD161) suggests a role for other cell sets, including natural killer (NK) cells and T lymphocytes in this lower disease state. The cell population expressing CD161 has been shown to be lower in SLE patients versus controls²⁹. This is intriguing as this receptor can mark the NK cells that respond to innate cytokines and so promote innate inflammation³⁰. Here again we see a contradiction between the promotion and reduction of the innate immune response.

While *CKAP4* was shown as a highly expressed hub gene in DA3, the protein product is most often reported to have a role in cancer, for example acting with RBP1 to induce autophagy in murine models of oral squamous cell carcinoma³¹. Autophagy can also play into the pathogenesis of SLE in a number of ways. Dysregulated autophagy can affect the regulation of T and B cell populations³², and increased autophagy can promote the NF- κ B pathway response³³. Through its interaction with ER-resident proteins, *CKAP4* also has the potential to regulate or reflect the current state of cellular immune signalling¹⁵. For the individuals studied here, increased levels of *CKAP4* may not be driving disease, but the finding opens a potential line of anti-*CKAP4* antibody drug development for SLE patients; an avenue previously only promoted for cancer treatment³⁴. Another DA3 hub gene, *OTOF*, is an interferon inducible gene, and has been recognised by others as a marker for SLE disease flares³⁵. This is in keeping with the finding of *OTOF* in the C3 and C5 clusters, but not in C4. Recently it was suggested that through interaction with melatonin, *OTOF* may have a role in proteasome inhibition³⁶, and so

could function in the downstream signal transduction pathway of NF- κ B³⁷. While that study was focused on neuronal survival driven by melatonin ubiquitin proteasome system inhibition, a protective anti-inflammatory role of melatonin in SLE pathogenesis has been reported previously^{38,39}. Gene networks acting through the full-crum of *OTOF* may help to explain this action, and suggests that further investigation of melatonin treatment in SLE flare could be warranted.

The current analysis aimed to explore the different networks that underlie pSLE disease states with the goal of developing a minimum set of rules that could discern disease states DA1 from DA3. It is worth mentioning that we did not aim to model the entire spectrum of pSLE disease activity, so we chose the objects that could optimally and clearly separate between DA1 and DA3 states and highlight their subgroups. This required the pruning the misclassified objects from the initial model. A logistic regression was used to reveal the probability of object removal (Supplementary Fig. S3 online, Supplementary Table S1). More specifically, this analysis showed that the enhanced model applies to a subset of SLE patients for whom there have been no prior treatments, do not have very low SLEDAI scores, or have not markedly long period of time between diagnosis and clinic visit. If the pathways are to be generalized, these factors must be accounted for prior to using the model in a clinical setting.

The enhanced model showed clearer sub-networks even though the gain in the accuracy was only 10%. While the networks generated here are based on a single gene expression set, multiple lines of evidence from previous SLE studies support their value; whether that be in classifying sub-cluster patient states or indicating possible treatments based on hub genes. It will be important to test the predictive, or replicative, ability of the gene networks to classify additional SLE patient sets. This includes generalisation to known patient subcategories, such as those with nephritis, and the wider adult SLE population. However, the permutation analysis conducted here suggests that this should be possible. We believe that machine learning approaches, such as the one demonstrated here, could aid disease understanding. This applies not only to SLE, but to any complex heterogeneous syndrome.

Methods

Figure 1, an overview of the analysis pipeline was generated with <http://www.lucidchart.com> resources.

Data and pre-processing. Publicly available whole blood transcriptome records (Illumina HT-12 V4 bead chip) and clinical metadata from 158 pSLE patients and 48 healthy controls were downloaded (NCBI GEO: GSE65391)⁹. All the procedures were performed in accordance with the relevant guidelines and regulations. From that existing data set, the values corresponding to DA1 (SLEDAI 0-2), DA3 (SLEDAI >7) and control visits were extracted. In this analysis, the transcriptome generated per visit to the clinic, and not per patient lifetime, was considered. As such, an individual may be represented in the analysis multiple times (between 1 and 15 times) if their disease status at the time was classified as DA1 or DA3 (Supplementary Fig. S12 online). For expression data, gene loci represented by more than one probe were combined and averaged, before each gene locus was log transformed. Following a linear mixed model used to identify biological and technical confounders (Variance Partition R package⁴⁰), no biological cofounders were identified. Batch effects were identified (Variance Partition R package⁴¹) and corrected (SVA R package⁴⁰). The batch effects identified here were limited to the reported batch replicates from the original metadata (batch 1 and 2) and not found for other phenotypes (Supplementary Fig. S13 online).

Machine learning rule-based modelling to obtain explainable classifiers for DA state. For methodological context, we applied an interpretable learning method based on rough sets that offers classification transparency^{11,12}. Given data in the form of a decision table, where rows represent observations and columns are features with the last column being the outcome or decision, rough set algorithms find all minimal subsets of features that preserve discernibility between the outcomes for the observations. These subsets of features are called reducts, and are used to generate IF-THEN rules by overlaying them on the observations. An IF-THEN rule consists of the condition part, the IF-part, often called the left-hand side, and the THEN part is the decision or outcome and often called the right-hand side of the rule. The elements of the IF-part are called descriptors, and are in the form of pairs, feature and its value. To aid interpretation, the rules generated by the model were visualized as RNs, where the nodes are descriptors. For every pair of descriptors in a rule of the RBM, an edge connecting the corresponding nodes is added to the network.

First, expression values were subject to data discretisation, since R. ROSETTA¹⁴ generates rules for that data form. For each gene, the control data expression mean (μ) and standard deviation (σ) were calculated, and then all DA data for that gene projected onto this threshold frame and discretised ($\text{Low} \leq \mu - 2\sigma$; $\mu - 2\sigma < \text{Medium} < \mu + 2\sigma$; $\text{High} \geq \mu + 2\sigma$; and coded by numeric values 1, 2, 3, respectively).

To generate the initial model, data was first collected into a decision table where unique visit identifiers were the objects and put in rows ($n = 629$), while genes ($n = 33,006$) were variables and constituted columns. The objects were labelled with disease activity, DA1 or DA3, accordingly. Next, Monte Carlo Feature selection (MCFS) algorithm¹⁴ was applied to obtain a ranked list of informative features with respect to classifying the objects. A significance cut-off for selecting features from the ranked list was obtained by a permutation test ($p\text{-value} \leq 0.05$). Feature boosting was applied to select the optimal number of features to build the model and then the rule model was visualized with the VisuNet R package⁴².

The initial rule-based model defined above was used as a base to further improve classification. Data (DA1 or DA3 visits) that did not match the left-hand side of any significant rules in the previous model were removed ($p\text{-value} < 0.05$). The MCFS¹⁵ process was then repeated after object removal. Prior to building the enhanced rule-based model, iterative computational rounds were performed (Feature boosting in Fig. 1) in order to select the optimal number of features for building the final predictive model. The significant features from MCFS output were incrementally added to build several rule-based models. The selected features that were used to build the

model with the best overall accuracy where chosen for building the final enhanced model using R.ROSETTA¹⁴ and then visualized using VisuNet⁴².

In order to identify patient subgroups, a matrix was constructed with maintained observations (visits) as rows and rules as columns. All cells for the observations that supported a rule were assigned 1, otherwise 0. Hierarchical clustering based on binary distance as the distance function was applied on this matrix.

Correlating clusters to clinical and phenotypic data. Available metadata, including continuous and categorical clinical values (Supplementary Table S3), were accessed¹⁰. For continuous variables, a one-way ANOVA following a post-hoc Tukey HSD test was used to compute significance. A Fisher's exact test was used for the assessment of categorical variables to sub-clusters.

Correlating rules associated with clusters to clinical and phenotypic data. Empirical values were used to determine the minimal threshold for rule membership to clusters. Rules were considered associated with a cluster if they had a support set matching at least 10% of the cluster's support set (i.e., observations associated with a cluster; Supplementary Fig. S14 online). The association between a cluster's supported rules and clinical phenotypes was assessed by contrasting phenotype values for supported samples of each rule versus the non-supported samples (categorical variables, non-parametric Wilcoxon test; binary variables, Fisher's exact test). Supplementary Fig. S15 online illustrates this process.

Model validation. The decision label (DA1 or DA3) was permuted 1000 times and rule-based models were created for these random sets. A normal distribution was built for the model accuracies and an alpha of 0.05 and a 95% confidence interval used to determine the significance of the p-value. The mean, standard deviation and the standard error for the normal distribution were computed. The accuracy of the original model was compared to the mean μ and standard error σ . If the accuracy of the original model was smaller than $\mu - \sigma$ or greater than $\mu + \sigma$ then the p-value in this case was < 0.05 .

Gene enrichment analysis. Overrepresentation of gene sets belonging to each cluster and the gene sets belonging to rules in DA1 and DA3 was determined using the R package clusterProfiler⁴³. The background list was set as initial set of 33,006 available loci.

Received: 8 June 2021; Accepted: 13 April 2022

Published online: 06 May 2022

References

- Kamphuis, S. & Silverman, E. D. Prevalence and burden of pediatric-onset systemic lupus erythematosus. *Nat. Rev. Rheumatol.* **6**(9), 538–546. <https://doi.org/10.1038/nrrheum.2010.121> (2010).
- Tsokos, G. C. Systemic lupus erythematosus. *N. Engl. J. Med.* **365**(22), 2110–2121. <https://doi.org/10.1056/NEJMra1100359> (2011).
- Peter, H. & Schur, B. C. Derivation and validation of the systemic lupus international collaborating clinics classification criteria for systemic lupus erythematosus—Petri—2012—Arthritis & Rheumatism—Wiley Online Library (Accessed 2 Mar 2021) <https://doi.org/10.1002/art.34473>.
- Bombardier, C. *et al.* Derivation of the sledai. A disease activity index for lupus patients. *Arthritis Rheum.* **35**(6), 630–640. <https://doi.org/10.1002/art.1780350606> (1992).
- Wang, Y. F. *et al.* Identification of 38 novel loci for systemic lupus erythematosus and genetic heterogeneity between ancestral groups. *Nat. Commun.* **12**(1), 772. <https://doi.org/10.1038/s41467-021-21049-y> (2021).
- Li, Q. Z. *et al.* Interferon signature gene expression is correlated with autoantibody profiles in patients with incomplete lupus syndromes. *Clin. Exp. Immunol.* **159**(3), 281–291. <https://doi.org/10.1111/j.1365-2249.2009.04057.x> (2010).
- Kyogoku, C. *et al.* Cell-specific type I IFN signatures in autoimmunity and viral infection: What makes the difference?. *PLoS ONE* **8**(12), e83776. <https://doi.org/10.1371/journal.pone.0083776> (2014).
- Demirkaya, E., Sahin, S., Romano, M., Zhou, Q. & Aksentijevich, I. New horizons in the genetic etiology of systemic lupus erythematosus and lupus-like disease: Monogenic lupus and beyond. *J. Clin. Med.* **9**(3), 712. <https://doi.org/10.3390/jcm9030712> (2020).
- Marion, T. N. & Postlethwaite, A. E. Chance, genetics, and the heterogeneity of disease and pathogenesis in systemic lupus erythematosus. *Semin. Immunopathol.* **36**(5), 495–517. <https://doi.org/10.1007/s00281-014-0440-x> (2014).
- Banchereau, R. *et al.* Personalized immunomonitoring uncovers molecular networks that stratify lupus patients. *Cell* **165**(3), 551–565. <https://doi.org/10.1016/j.cell.2016.03.008> (2016).
- Komorowski, J. 6.02—Learning rule-based models—the rough set approach. In *Comprehensive Biomedical Physics* (ed. Brahmé, A.) 19–39 (Elsevier, 2014). <https://doi.org/10.1016/B978-0-444-53632-7.01102-3>.
- Skowron, A. & Dutta, S. Rough sets: Past, present, and future. *Nat. Comput.* **17**(4), 855–876. <https://doi.org/10.1007/s11047-018-9700-3> (2018).
- Pawlak, Z. Rough sets. *Int. J. Comput. Inf. Sci.* **11**(5), 341–356. <https://doi.org/10.1007/BF01001956> (1982).
- Garbulowski, M. *et al.* R.ROSETTA: An interpretable machine learning framework. *BMC Bioinform.* **22**(1), 110. <https://doi.org/10.1186/s12859-021-04049-z> (2021).
- Dramiński, M. *et al.* Monte Carlo feature selection for supervised classification. *Bioinformatics* **24**(1), 110–117. <https://doi.org/10.1093/bioinformatics/btm486> (2008).
- Gao, G., Zhu, C., Liu, E. & Nabi, I. R. Reticulon and CLIMP-63 regulate nanodomain organization of peripheral ER tubules. *PLoS Biol.* **17**(8), e3000355–e3000355. <https://doi.org/10.1371/journal.pbio.3000355> (2019).
- Xiahou, Z. *et al.* NMI and IFP35 serve as proinflammatory DAMPs during cellular infection and injury. *Nat Commun.* **8**, 1–11. <https://doi.org/10.1038/s41467-017-00930-9> (2017).
- Bradley, S. J., Suarez-Fueyo, A., Moss, D. R., Kyttaris, V. C. & Tsokos, G. C. T cell transcriptomes describe patient subtypes in systemic lupus erythematosus. *PLoS ONE* **10**(11), e0141171. <https://doi.org/10.1371/journal.pone.0141171> (2015).
- Harley, J. B. *et al.* Genome-wide association scan in women with systemic lupus erythematosus identifies susceptibility variants in ITGAM, PXX, KIAA1542 and other loci. *Nat. Genet.* **40**(2), 204–210. <https://doi.org/10.1038/ng.81> (2008).

20. Löfgren, S. E. *et al.* Genetic association of miRNA-146a with systemic lupus erythematosus in Europeans through decreased expression of the gene. *Genes Immun.* **13**(3), 268–274. <https://doi.org/10.1038/gene.2011.84> (2012).
21. Toro-Domínguez, D. *et al.* Stratification of systemic lupus erythematosus patients into three groups of disease activity progression according to longitudinal gene expression. *Arthritis Rheumatol.* **70**(12), 2025–2035. <https://doi.org/10.1002/art.40653> (2018).
22. Tesser, A. *et al.* Higher interferon score and normal complement levels may identify a distinct clinical subset in children with systemic lupus erythematosus. *Arthritis Res Ther.* **22**(1), 91. <https://doi.org/10.1186/s13075-020-02161-8> (2020).
23. Han, B. K. *et al.* Neutrophil and lymphocyte counts are associated with different immunopathological mechanisms in systemic lupus erythematosus. *Lupus Sci. Med.* **7**(1), e000382. <https://doi.org/10.1136/lupus-2020-000382> (2020).
24. Greene, C. S., Tan, J., Ung, M., Moore, J. H. & Cheng, C. Big data bioinformatics. *J. Cell Physiol.* **229**(12), 1896–1900. <https://doi.org/10.1002/jcp.24662> (2014).
25. Azodi, C. B., Tang, J. & Shiu, S. H. Opening the black box: Interpretable machine learning for geneticists. *Trends Genet.* **36**(6), 442–455. <https://doi.org/10.1016/j.tig.2020.03.005> (2020).
26. Glaab, E., Bacardit, J., Garibaldi, J. M. & Krasnogor, N. Using rule-based machine learning for candidate disease gene prioritization and sample classification of cancer gene expression data. *PLoS ONE* **7**(7), e39932. <https://doi.org/10.1371/journal.pone.0039932> (2012).
27. Bange, F. C. *et al.* IFP 35 is an interferon-induced leucine zipper protein that undergoes interferon-regulated cellular redistribution. *J. Biol. Chem.* **269**(2), 1091–1098 (1994).
28. Das, A., Dinh, P. X., Panda, D. & Pattnaik, A. K. Interferon-inducible protein IFI35 negatively regulates RIG-I antiviral signaling and supports vesicular stomatitis virus replication. *J. Virol.* **88**(6), 3103–3113. <https://doi.org/10.1128/JVI.03202-13> (2014).
29. Lin, Y. L. & Lin, S. C. Analysis of the CD161-expressing cell quantities and CD161 expression levels in peripheral blood natural killer and T cells of systemic lupus erythematosus patients. *Clin. Exp. Med.* **17**(1), 101–109. <https://doi.org/10.1007/s10238-015-0402-1> (2017).
30. Kurioka, A. *et al.* CD161 defines a functionally distinct subset of pro-inflammatory natural killer cells. *Front. Immunol.* **9**, 486. <https://doi.org/10.3389/fimmu.2018.00486> (2018).
31. Gao, L. *et al.* The RBP1-CKAP4 axis activates oncogenic autophagy and promotes cancer progression in oral squamous cell carcinoma. *Cell Death Dis.* **11**(6), 488. <https://doi.org/10.1038/s41419-020-2693-8> (2020).
32. Clarke, A. J. *et al.* Autophagy is activated in systemic lupus erythematosus and required for plasmablast development. *Ann. Rheum. Dis.* **74**(5), 912–920. <https://doi.org/10.1136/annrheumdis-2013-204343> (2015).
33. Zhong, Z. *et al.* NF- κ B restricts inflammasome activation via elimination of damaged mitochondria. *Cell* **164**(5), 896–910. <https://doi.org/10.1016/j.cell.2015.12.057> (2016).
34. Kikuchi, A., Fumoto, K. & Kimura, H. The Dickkopf1-cytoskeleton-associated protein 4 axis creates a novel signalling pathway and may represent a molecular target for cancer therapy. *Br. J. Pharmacol.* **174**(24), 4651–4665. <https://doi.org/10.1111/bph.13863> (2017).
35. Mackay, M. *et al.* Molecular signatures in systemic lupus erythematosus: Distinction between disease flare and infection. *Lupus Sci. Amp Med.* **3**(1), e000159. <https://doi.org/10.1136/lupus-2016-000159> (2016).
36. Yalcin, E. *et al.* Evidence that melatonin downregulates Nedd4-1 E3 ligase and its role in cellular survival. *Toxicol. Appl. Pharmacol.* **379**, 114686. <https://doi.org/10.1016/j.taap.2019.114686> (2019).
37. Bruck, R. *et al.* Melatonin inhibits nuclear factor kappa B activation and oxidative stress and protects against thioacetamide induced liver damage in rats. *J. Hepatol.* **40**(1), 86–93. [https://doi.org/10.1016/S0168-8278\(03\)00504-X](https://doi.org/10.1016/S0168-8278(03)00504-X) (2004).
38. Huang, H., Liu, X. & Chen, D. Corrigendum to “Melatonin prevents endothelial dysfunction in SLE by activating the nuclear receptor retinoic acid-related orphan receptor- α ” [Int. Immunopharmacol. **83** (2020) 106365]. *Int. Immunopharmacol.* **2020**(86), 106817. <https://doi.org/10.1016/j.intimp.2020.106817> (2020).
39. Bonomini, F., Dos Santos, M., Veronese, F. V. & Rezzani, R. NLRP3 inflammasome modulation by melatonin supplementation in chronic pristane-induced lupus nephritis. *Int. J. Mol. Sci.* **20**(14), 3466. <https://doi.org/10.3390/ijms20143466> (2019).
40. Hoffman, G. E. & Schadt, E. E. variancePartition: Interpreting drivers of variation in complex gene expression studies. *BMC Bioinformatics* **17**(1), 483. <https://doi.org/10.1186/s12859-016-1323-z> (2016).
41. Leek, J. T., Johnson, W. E., Parker, H. S., Jaffe, A. E. & Storey, J. D. The sva package for removing batch effects and other unwanted variation in high-throughput experiments. *Bioinform. Oxf. Engl.* **28**(6), 882–883. <https://doi.org/10.1093/bioinformatics/bts034> (2012).
42. Lab KB. *Komorowskilab/VisuNet*. (2020). <https://github.com/komorowskilab/VisuNet>. (Accessed 2 Mar 2021).
43. Yu, G., Wang, L. G., Han, Y. & He, Q. Y. clusterProfiler: An R package for comparing biological themes among gene clusters. *OMICS J. Integr. Biol.* **16**(5), 284–287. <https://doi.org/10.1089/omi.2011.0118> (2012).

Acknowledgements

We would like to thank Mateusz Garbulowski for providing the rule networks figure under “Disease subtype discovery” in Fig. 1 in addition to R. ROSETTA and VisuNet logos.

Author contributions

J.K. conceptualized the experiment with design input from C.F.B., J.R.S.M., and S.A.Y. J.R.S.M. and S.A.Y. wrote the manuscript. P.S. created the initial model and A.A. enhanced R scripts. K.D. refined the permutation test, implemented the use of binary distance for hierarchal clustering and enhanced the pruning steps for the enhanced model. C.F.B. designed the cluster-to-phenotype enrichment analysis. S.A.Y. performed all remaining modelling and intersection with phenotypic data. C.F.B., J.R.S.M. and S.A.Y. interpreted the results. J.R.S.M. developed the discussion. L.H. refined manuscript text. S.A.Y. generated all figures and tables with support from J.R.S.M for Figs. 3 and 4, and all Supplementary Tables. All authors reviewed the manuscript.

Funding

Open access funding provided by Uppsala University. J.R.S.M. was supported by the Swedish Research Council, FORMAS (221-2012-1531). Open Access funding provided by Uppsala University. The salary of S.A.Y. was supported by a grant from the Knut and Alice Wallenberg Foundation held by L.H. (KAW 2013-0159). J.K. was funded in part by the Polish Academy of Sciences, Institute of Computer Science. J.K. and C.F.B. were funded in part by The University of Washington, Seattle, The National Institute of Allergy and Infectious Diseases, Division of AIDS, National Institutes of Health (ABL Contract No. HHSN272201700010I) and by the eSSence program. The funders had no role in study design, data collection and analysis, decision to publish, or preparation of the manuscript.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary Information The online version contains supplementary material available at <https://doi.org/10.1038/s41598-022-10853-1>.

Correspondence and requests for materials should be addressed to S.A.Y. or J.K.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.



Open Access This article is licensed under a Creative Commons Attribution 4.0 International License, which permits use, sharing, adaptation, distribution and reproduction in any medium or format, as long as you give appropriate credit to the original author(s) and the source, provide a link to the Creative Commons licence, and indicate if changes were made. The images or other third party material in this article are included in the article's Creative Commons licence, unless indicated otherwise in a credit line to the material. If material is not included in the article's Creative Commons licence and your intended use is not permitted by statutory regulation or exceeds the permitted use, you will need to obtain permission directly from the copyright holder. To view a copy of this licence, visit <http://creativecommons.org/licenses/by/4.0/>.

© The Author(s) 2022