Computational and Structural Biotechnology Journal 19 (2021) 4751–4758

COMPUTATIONAL AND STRUCTURAL BIOTECHNOLOGY JOURNAL

journal homepage: www.elsevier.com/locate/csbj

Regulatory network analysis defines unique drug mechanisms of action and facilitates patient-drug matching in alopecia areata clinical trials

James C. Chen *, Zhenpeng Dai, Angela M. Christiano *

Department of Dermatology, Columbia University Medical Center, United States

ARTICLE INFO

Article history:
Received 29 April 2021
Received in revised form 16 August 2021
Accepted 16 August 2021
Available online 19 August 2021

Keywords:
Systems biology
Alopecia areata
Artificial intelligence
Drug prediction
Personalized medicine

ABSTRACT

Not all therapeutics are created equal in regards to individual patients. The problem of identifying which compound will work best with which patient is a significant burden across all disease contexts. In the context of autoimmune diseases such as alopecia areata, several formulations of JAK/STAT inhibitors have demonstrated efficacy in clinical trials. All of these compounds demonstrate different rates of response, and here we observed that this coincided with different molecular effects on patients undergoing treatment. Using these data, we have developed a computational model that is capable of predicting which patient-drug pairs have the highest likelihood of response. We achieved this by integrating inferred mechanism of action data and gene regulatory networks derived from an independent patient cohort with baseline patient data prior to beginning treatment.

1. Introduction

Imagine that a patient is newly diagnosed with alopecia areata (AA). The next question that arises is a crucial one: which treatment should the clinician begin with? Several promising treatments have emerged for the disease, yet none of them are trivial in either expense, or the risk of side effects, and there is always a risk that a treatment is ineffective. This problem extends to the many diseases for which many treatment options are available – who should get what?

JAK/STAT-class inhibitors have emerged as a promising therapeutic option for the treatment of several autoimmune diseases, and several formulations have undergone clinical trial testing specifically for AA. Our group and others have evaluated the efficacy of compounds such as tofacitinib (tofa), a pan-JAK inhibitor, and ruxolitinib (ruxo), a JAK1/2 inhibitor, in the treatment of AA [1–3] in open-label clinical trials. In addition, we have tested compounds such as abatacept (CTLA4-Ig) and standard-of-care treatments such as intralesional triamcinolone injections (ILTAC) [4].

During the course of these trials, we observed a variable rate of non-response across each study, ranging from 20 to 40% of patients as defined by <50% resolution of SALT score [5]. As each trial was conducted independently, each treated patient was enrolled for a specific clinical trial without consideration for any of the other available compounds.

However, now that we have several validated and promising options, each with variable response rates, the question we now seek to address with this data is whether or not it is possible to sensibly decide which patient gets which treatment before they begin any treatment at all? While numerous studies have provided extensive molecular quantification of varying immune processes and "immune privilege" mechanisms governing AA [4,6–9], there has been little success in translating these observations to clinical predictors to specific therapeutics. While many of these models provide promising accounts of the underlying pathology, they do not directly consider the effects of drug treatments such as JAK inhibitors, which are not specifically designed to target these biomarker panels, but rather intersect with them in some way. Our inclusion of biopsy sampling throughout the clinical trials allowed us to measure the molecular response to treatment through RNA-seq profiling alongside clinical outcome. The molecular effects of each treatment course were individually related back to the clinical outcomes via the (ALopecia Areata Disease INdex) ALADIN score in the individual trials [4]. With these trials now complete, we can use the data to drive more advanced computational modeling of drug mechanism of response (MoR), as these trials have provided us with a time course of the molecular effects of each compound in an *in vivo* setting, through RNAseq analysis of the patient scalp biopsies.

* Corresponding authors.
E-mail addresses: jameschen0413@gmail.com (J.C. Chen), amc65@cumc.columbia.edu (A.M. Christiano).

https://doi.org/10.1016/j.csbj.2021.08.026
2001-0370/Published by Elsevier B.V. on behalf of Research Network of Computational and Structural Biotechnology.

Here, we integrated the RNAseq data for these trials into a single framework using reverse-engineered regulatory networks [9,10]. These networks describe the theoretical transcriptional regulatory logic of target tissues (in this instance, the scalp skin and hair follicle) by inferring transcriptional interactions directly from transcriptome profiling. The algorithm allows us to describe large pathogenic molecular signatures as a smaller core of transcriptional regulators that are required to maintain these signatures, which we refer to as candidate master regulators (MRs). These MRs were then compared against molecular models for each compound's MoR to define the effects of each drug.

Due to the limitation of leveraging already-completed clinical trials specifically in the context of AA, e.g. retrospective, relatively limited sample sizes, we designed a specific framework of prediction concerned primarily with predicting non-response rates in specific compound-patient matchups, and aimed to create a computational model that could predict non-response rates of individual patients to each compound using only the initial, pre-treatment biopsies. This study focused specifically on rank-ordering these three compounds based on the predicted non-response likelihood. This has the caveats that 1) it was not designed to discern which of the compounds would certainly work (i.e. just because the algorithm did not pick the drug as the best does not imply that the drug would fail), and 2) it does not have predictive coverage of drugs that were not included in the analysis (but they can be included if the corresponding data becomes available). The aim of the study was to provide highest-likelihood predictions to minimize non-response rates.

This is a crucial requirement for clinical application, since the goal is to be able to aid in therapeutic decisions before putting patients on treatments. This computational framework provides the foundation for bioinformatic patient-matching of individual patients to the treatments with the highest probability of positive response. Fig. 1 provides a conceptual workflow of the algorithm.

2. Results

2.1. Naïve differential expression analysis of four independent AA clinical trials defines distinct drug response programs

Initially, we collated the patient data for the four clinical trials (ruxo, tofa ILTAC, and abatacept) as four independent groups and performed a standardized, naïve differential expression analysis for each cohort as a baseline (see methods). The molecular mechanism of response (MoR) signature was defined as the genes differentially expressed between each treatment groups' endpoint treatment biopsy and its corresponding pre-treatment biopsy. For reference, the distribution of differential expression and statistical significance are presented in Fig. 1 as volcano plots (Fig. 2 A-D, top panels).

For each group, these MoR signatures were used in an unsupervised hierarchical clustering analysis with unaffected controls spiked in to assess the strength of MoR gene clustering (Fig. 1 A-D heatmaps). It should be noted that these gene signatures are not explicitly informed by the ALADIN signature, and exist largely orthogonally to ALADIN. Tofa and ruxo were associated with the most robust molecular response, defined both by the separation of pre- and post-treatment samples and by the cluster strength of the MoR gene sets. ILTAC and abatacept overall associated with fewer differentially expressed genes and reduced clustering quality. Both of these compounds also associated with the lowest overall clinical response rates in the clinical trials.

2.2. Direct comparison of four compounds reveals orthogonal mechanisms of action in both human clinical trials and mouse models

Using the response signatures above, we can compare the gene sets associated to each course of treatment and assess their overall similarity. Fig. 2E provides a representation of the relative overlap of genes that are differentially expressed in association with each completed course of treatment. The accompanying table presents the p-value of the corresponding left-tailed Fisher's Exact test for each pair of gene sets – a statistically significant p-value for the pair indicates that the two gene sets are statistically non-overlapping and distinct from each other. From this analysis, the MoR sets of all four compounds were statistically independent of each other, including the two JAK inhibitors, tofacitinib and ruxolitinib, suggesting that all four of the tested compounds have separate, traceable molecular MoRs.

These results were supported by mouse models in which we treated mice with AA using tofacitinib, ruxolitnib, and compounds targeting specific JAK pathway genes. Specifically, the unique MoRs of ruxo and tofa observed in human cohorts was recapitulated in the mouse models (supplemental Tables 1–4 for relevant data) [11]. While tofacitnib and ruxolitinib exhibit molecular perturbation behaviors consistent with combinations of JAK1 and JAK3 targeted compounds, the topography of their respective graphs indicate synergistic effects greater than the simple sum of the two compounds, expressed as the significant expansion of over-represented genes in the MoR signatures.

These results suggest that the molecular downstream effects of treatment in the *in vivo* setting can vary drastically beyond what may be suggested by the class of compound, e.g. JAK1 vs JAK3 vs JAK1/3. These molecular complexities produce a system that is difficult to model using consensus-based differential expression, but also allow for the potential for patient-matching optimization if an integrative representation of MoR can be modeled.

2.3. Master regulator analysis of the gene networks of treatment response in these clinical trials reveals distinct modules of unique mechanism of response

Having defined these orthogonal MoR signatures, we performed a master regulator analysis to identify the core, minimal set of transcriptional regulators required to propagate and recapitulate this signature using our existing, validated skin regulatory network [9]. This will eventually have predictive capability, but we initially must define a set of MRs for each drug. Each master regulator analysis was conducted independently for each MoR signature, and the networks were integrated and visualized in Fig. 2. Of the four compounds tested, only abatacept had an MoR signature that failed to map to any individual regulatory hub with a coverage of > 70%, which is considered a heuristic measure of confidence [9,12–14]. This could suggest that the MoR of abatacept does not significantly propagate through transcriptional machinery or that it may not significantly affect the molecular steady state of the scalp tissue. For biomarker and clinical prediction purposes, we then iterated through this combined network and eliminated the edges shared between multiple treatments to create a treatment network module reflecting only unique regulatory interactions associated with each MoR (Fig. 2F,G). At the chosen statistical significance level (see methods), the MoR of these four compounds can be described and differentiated through a combination of six sets of MRs. While these intersected edges may be biologically relevant, for prediction purposes we focused on maximizing the ability to differentiate between the drugs.

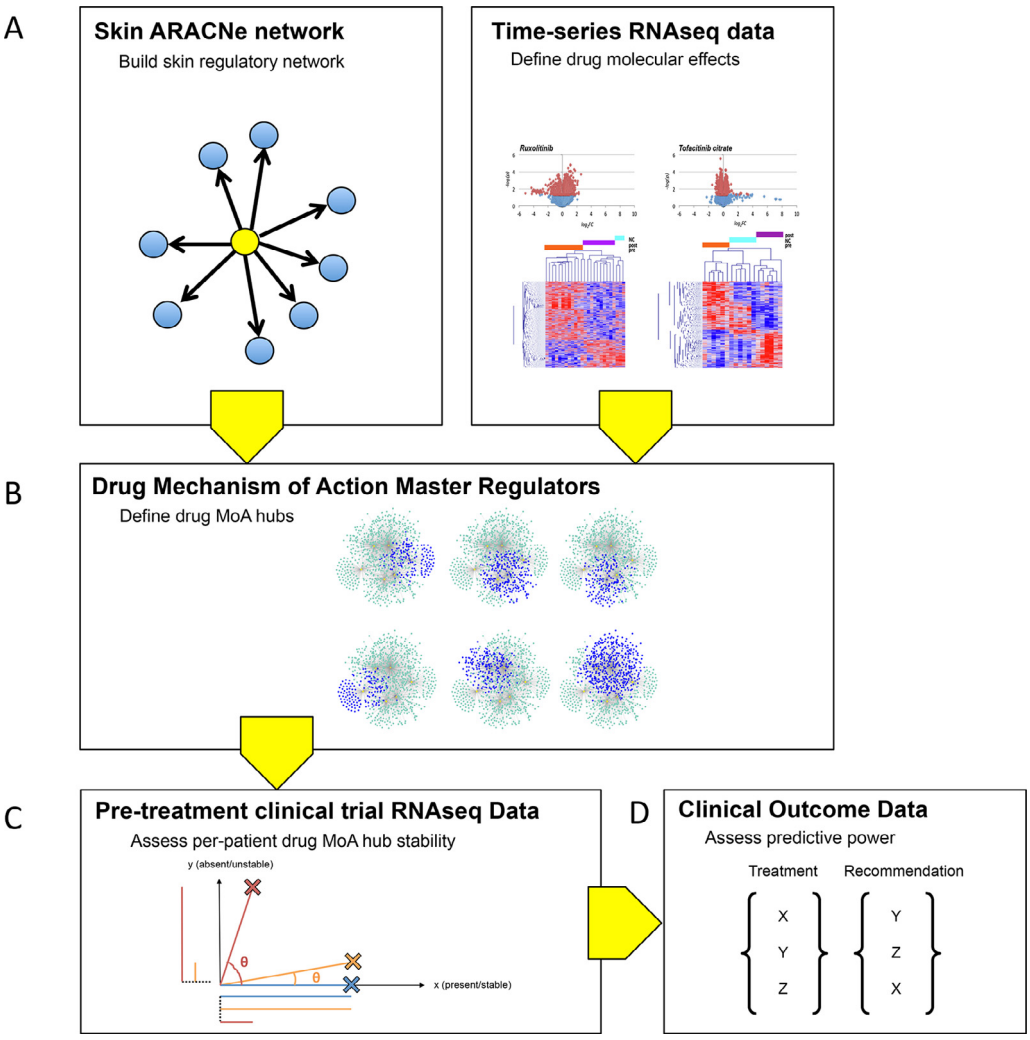

Fig. 1. A schematic flow diagram of the progression of this algorithm. (A) An independent skin regulatory network is integrated with drug mechanism of response data inferred from a clinical trial time-series. (B) This data is used to map the master regulators that define molecular response to each compound. These master regulators uniquely describe the molecular hubs associated to each compound. (C) on a per-patient basis, we assess the transcriptional state of each hub in relation to the theoretical master regulator hub provided by our network. The concordance is quantified for each hub in each patient. (D) Treatments for each patient are ranked based on the strength of concordance between the observed molecular state and the theoretical network hub. Patient response is then unblinded and accuracy in outcome is measured.

2.4. Molecular response programs can be used to predict non-responder rates in patient cohorts prior to treatment and to rank responsiveness to treatment

Using the master regulator hubs generated from Fig. 2G, we developed a model to measure the similarity of an individual patient's response hub its corresponding computationally predicted model (see Methods for details). The method essentially identifies patients whose expression an MR-predicted target falls outside the reference set normal range. These outlier genes are subsequently marked as network edges that do not conform to the computationally inferred network. The overall prediction of patient response is then defined as the relative concordance between the computationally inferred network and the number of these dysregulated edges in an individual patient by computing the total hub co-information minus the information of the dysregulated edges.

When performing this analysis across multiple patients, we observed that each patient had unique set of deviations from the inferred hub (Fig. 3A,B). This cartoon, and the associated figure, shows an abstraction of this process (the full details are available in Methods). Given a specific regulatory hub, e.g. tofacitinib, each of three patients are assessed for genes in the hub that exist at the appropriate steady state (present) or not (absent). These edges are then weighted by their network-predicted mutual information into a bivariate vector that measures the overall relationship between total present and absent information.

The angle, theta, of this vector therefore represents the deviation of this individual patient's regulatory hub from the prototypical, inferred ARACNe hub (see Methods for mathematical proof of metrics). Here, a greater theta demonstrates less concordance between the master regulator state of an individual patient with the theoretical MR hub, and we interpret this as a lower likelihood of response to the associated treatment.

When done for an individual patient's pre-treatment biopsies, and using the MoR hubs of multiple compounds, it becomes possible to rank each hub according to its concordance with the ARACNe-predicted hub (Fig. 3C). Each compound can be mapped to a theta value and ranked, lowest to highest corresponding to best- and worst choice, respectively, for each patient. Note that the current implementation of this method does not provide information as to whether each individual treatment will or will not work, rather, it provides the best-case estimate of the best candidate drug based on the patient's pre-treatment molecular signature.

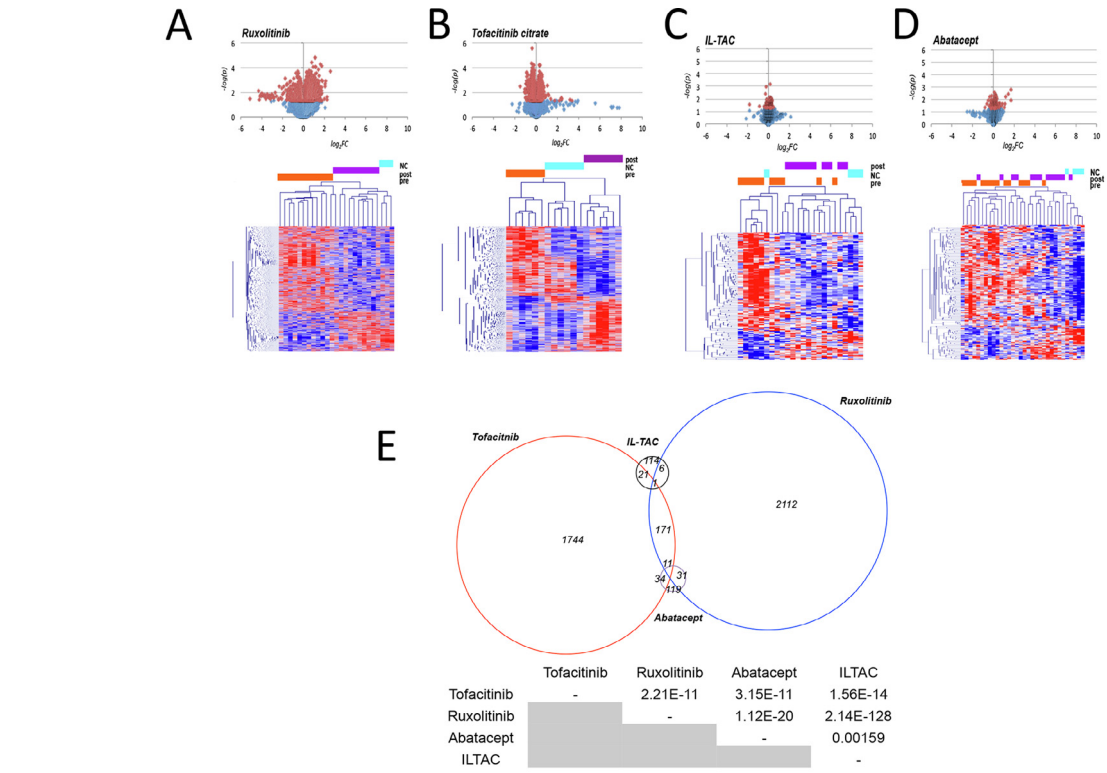

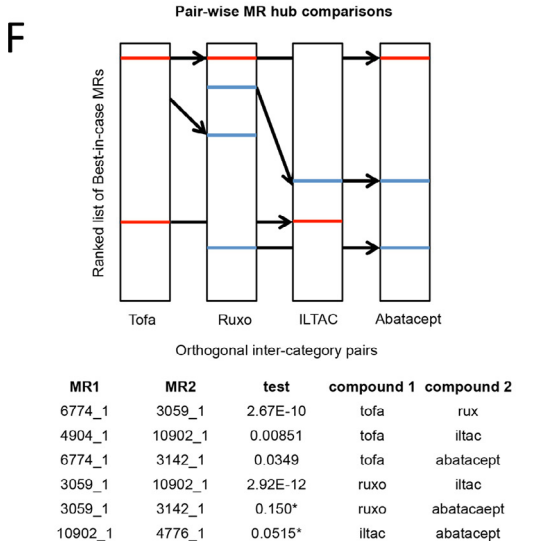

MR1	MR2	test	compound 1	compound 2
6774_1	3059_1	2.67E-10	tofa	rux
4904_1	10902_1	0.00851	tofa	iltac
6774_1	3142_1	0.0349	tofa	abatacept
3059_1	10902_1	2.92E-12	ruxo	iltac
3059_1	3142_1	0.150*	ruxo	abatacaept
10902_1	4776_1	0.0515*	iltac	abatacept

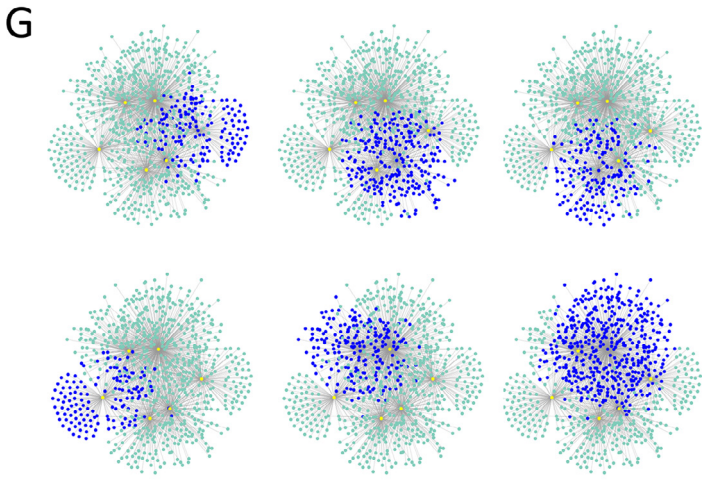

Fig. 2. Biomarker panels associated with prolonged treatment of (A) ruxolitinib, (B) tofacitinib citrate, (C) intralesional triamcinolone, and (D) abatacept were generated using RNAseq data generated from scalp biospies taken throughout the clinical trial. Each compound had variable response in terms of both the number of genes (volcano plots) and the robustness of the response (heatmaps) in separating pre-treatment (heatmap, orange barcodes) from post-treatment (heatmap, purple barcodes) and unaffected controls (heatmap, teal barcodes). (E) Venn diagram showing the overlap of the biomarker panels associated with each treatment course and accompanying pairwise overlap statistics. Sufficient statistical evidence exists that all biosignature panels are non-overlapping $p < 0.05$ corrected. (F) These signatures are integrated with a skin regulatory network using an edge-crawl method to identify master regulators with the maximal unique regulatory logic associated with each drug, allowing for (G) the mapping of distinct master regulators (yellow nodes) to each compound for outcome prediction. The MoRs of the four tested compounds can be covered through the combination of six candidate MR hubs. Dark blue nodes indicate unique MR gene hubs that differentiate the compounds, the overall cloud shows the total molecular hub describing all response genes in the study. (For interpretation of the references to colour in this figure legend, the reader is referred to the web version of this article.)

2.5. Ranking treatments per-patient by theta concordance associates with observed response rates

To assess the accuracy and application of this method, we then unblinded the patient treatment and response rates, and compared the best-case prescription by the algorithm to the treatment the patient actually received, and whether or not the patient exhibited a response at the end of a 24-week treatment cycle, defined as either at least 20% response in SALT score or by whether or not they were predicted responders by ALADIN (Fig. 3D).

If a patient was treated by a compound that the algorithm did not mark as best-in-case, and that patient failed to respond, this event was interpreted as a correct identification of non-response (the algorithm suggested a different treatment, and the patient did not respond on their prescribed treatment). If the patient was treated with a compound that the algorithm predicted was best,

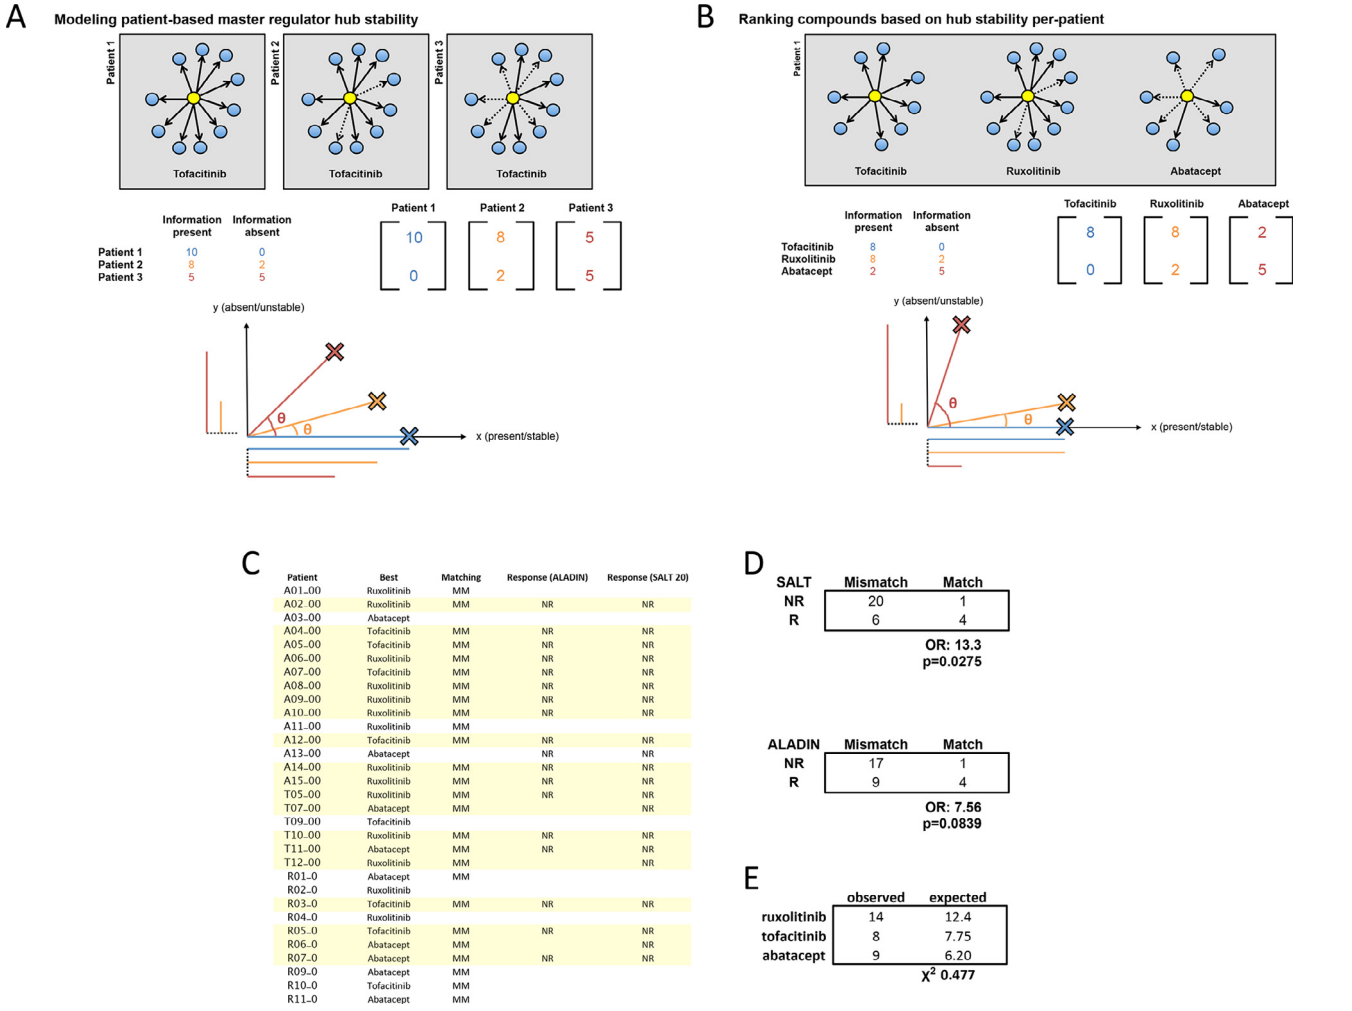

Fig. 3. (A) Schematic representation of measuring master regulator hub stability of a tofacitinib master regulator in three different patients. The same theoretical regulatory hub of a master regulator is compared against the molecular state of three different patients, all of whom have varying levels of concordance with it. Using mutual information this level of concordance can be vectorized into quantified concordant and discordant edges. Concordance overall is measured by angle theta of this vector, with lowest values being most concordant with the model, and therefore most likely to elicit clinical response, e.g. Patient 1 would be the most likely of the three to respond to tofacitinib. (B) The same as A, except now measuring three different compounds in an individual patient for ranking, e.g. for this Patient 1, tofacitinib would be the most likely to elicit a response, and abatacept the least likely. (C) taking each individual patient's pre-treatment biopsy and applying this logic provides a ranking for best treatment of the options available. Patients are then unblinded and their given treatment and response rate are compared with the algorithm's prediction. (D) the algorithm has significant association between first-pick and patient response, and mismatch with patient non-response using two different clinical outcome measures. (E) the algorithm assigns available compounds in proportions concurrent with their expected response rates.

and the patient responded, this was interpreted as a correct prediction and match, etc. Based on these interpretations, the categorical tabulation of prediction vs response rates are provided for both scoring definitions (ALADIN and SALT), see supplemental Tables 5–9 for related information. Both were statistically associated at alpha < 0.05 corrected.

Furthermore, we tabulated the algorithm's overall assignment rates of each of the compounds, e.g. how often/for how many patients did the algorithm suggest ruxolitib, and compared those rates to the expected response rates in the cohort based on the non-responder rates in each individual clinical trial. This was to assess the potential "bias" of the algorithm in picking "ideal" compounds – if the algorithm categorically assigned ruxolitinib to all patients (perhaps because it has the overall highest response rate of the trials conducted), it loses its value as a predictive matching algorithm. However, the distribution of the three compounds was concordant with the expected rates given overall observed clinical response rates by chi-squared test of independence ($x^2 = 0.477$, not significant, Fig. 3E). This result suggests that the algorithm's relative rates of assigning each compound are concordant with overall observed response rates associated with each compound.

2.6. Edge-driven patient-matching has significant positive and negative predictive value pre-treatment

Based on these results, we proceeded with an assessment of the algorithm's positive and negative predictive values based on posterior probabilities. The main findings are derived and detailed below in Fig. 4 (details in Methods). In summary, in this study the algorithm has a 78% (Fig. 4A) probability of correctly predicting patient response in a pretreatment biopsy, i.e., when a patient was given the treatment predicted by the algorithm to be the best fit for the patient, the patient did respond in 78% of cases. This is compared to the overall response rates of ~ 20–70% for each of the individual clinical trials, for a LR improvement of 1.3–3.9. Conversely, the algorithm had a negative predictive value overall of 65%, i.e., when a patient was given a compound that was not the best-in-case prediction by the algorithm, 65% of the patients failed to

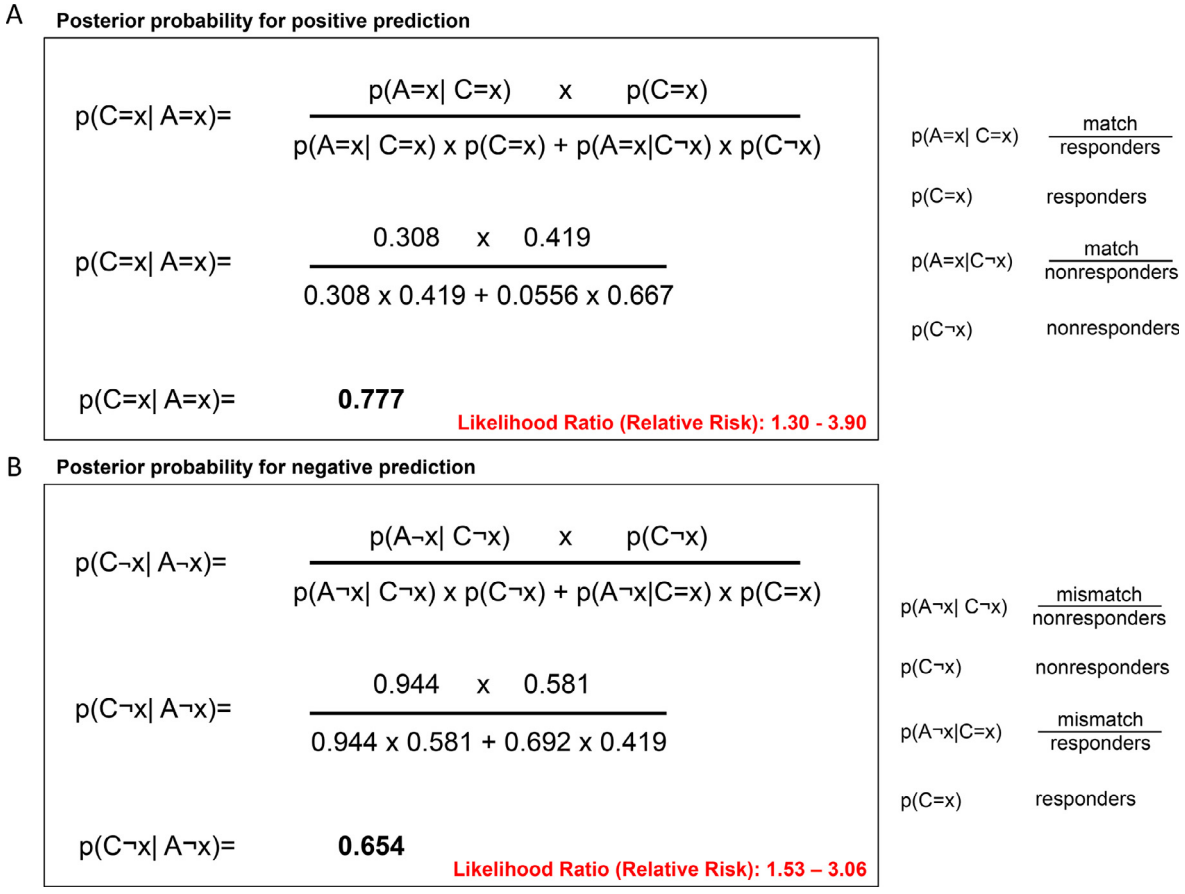

Fig. 4. (A) an assessment of the positive predictive value of the algorithm. Using the unblinded data, the posterior probability of patient response to a treatment given that the algorithm assessed that treatment as the best-case for that patient is 77.7%, with a relative risk improvement of 1.3–3.9, depending on the base response rate of the particular compound. (B) the assessment of negative predictive value based on posterior probability. The probability of a patient failing to respond to a treatment given that the algorithm did not select that treatment as the best-case option is 65.4%, with a relative risk improvement of 1.53–3.06 depending on the base response rate of the particular drug.

respond to the treatment, with an overall LR improvement of 1.53–3.6 (Fig. 4B). The lower bound of improvement (1.53) is due to comparing the overall matching rate to ruxolitinib, which overall had the highest response rate (and consequently the lowest non-response rate), i.e. ruxolitinib has the highest response rate across the individual trials, so comparing global probability assessments against it yields lower improvements in likelihood simply because the base response rate is higher.

3. Discussion

This study provides molecular quantification supporting the hypothesis that these four tested compounds have drastically different molecular fingerprints in scalp biopsies of AA patients. Regardless of the convergence of their molecular pathology, as measured by the ALADIN score, the exact patient response as measured by gene expression changes suggest that the biological pathways through which the suppression of pathology is achieved is unique to each compound. From even the most basic level of differential gene expression, statistical analysis can be conducted to show that the molecular effects of each of these compounds is statistically differentiable from each other, and distinct in signature. While a limited number of genes do fall within the intersections of each MoR signature, on the whole each compound elicits a unique molecular response in a treatment cohort that can be used to infer what compound a patient was subject to with no a priori information. In essence, patients can arrive at a remitted disease through multiple different molecular paths. This finding in itself warrants further investigation, as there has been little research devoted specifically to tracking the molecular effects of compounds on the patients in clinical trials.

One of the more striking findings of this study was the extent of the molecular divergence in MoR of different formulations of JAK inhibitors. The molecular MoRs associated with tofacitinib (pan-JAK) and ruxolitinib (JAK1,2) in the context of human clinical trials is statistically robust and orthogonal, and are as disparate from each other as they are from ILTAC and abatacept. In fact, all for compounds tested in human clinical trials had statistically distinct MoRs. These observations were corroborated in controlled, pre-clinical animal model studies of targeted JAK1, -2, and -3 inhibitors alongside ruxolitinib and tofacitinib.

Individual JAK compounds have statistically separable molecular effects, and the combinatorial effects of pan- and mult-JAK inhibitors are not a simple additive or linear progression. There are genes within the intersections of these MoRs, but while these individual genes may be biologically interesting, the global view of the MoR suggests the importance of regarding these compounds as distinct and non-equivalent treatments. Comparing the multi-targeting tofacitinib and ruxolitinib compounds to the individual JAK inhibitors directly shows a perhaps-surprising lack of overlap in gene signatures at all. Rather than each pan-JAK inhibitor being a superset of JAK1 and JAK3 inhibitor effects (or of all three), there is instead no statistical evidence of overlap of any pairwise comparison (no comparisons pass a $p < 0.05$ FDR threshold).

The global differences in these molecular MoRs beget the possibility of dynamic matching of individual patients to specific compounds in order to minimize non-response rates. These differences also highlight the importance of not implicitly regarding different compounds targeting the same pathway as interchangeable and, as a corollary, the potential dangers of doing so: these data show that these various JAK treatments are not in any way interchangeable as a function of the total molecular perturbation associated with patient response.

Here, we investigated the hypothesis that differences in response rates could be predicted on a per-patient basis by using the pre-treatment, steady-state molecular level of key master regulators associated with the drugs' MoRs. Published research from other groups [15,16] established the feasibility of using computational methods such as consensus-based master regulator analysis to model drug mechanism of response. Such methods are significant in the scope of both designing novel drugs and predicting response because it gives us a quantitative method to assess and predict how large molecular pathologies progress through treatment and time. Here, we leveraged our regulatory networks to collapse these MoRs into a core set of master regulators. For each compound pair, we mapped and identified the best non-overlapping hubs of molecular MoR for bioinformatic predictions. Using these hubs, we were able to develop a computational framework to estimate how well an individual patient would respond to each available compound based on their pre-treatment molecular state, and demonstrated that these predictions could be related back to the response rate predicted by the ALADIN signature and in clinical recovery metrics via the SALT score.

Despite the predictive power of the algorithm, it is important to note that its current implementation is not designed to assess whether or not specifc compounds tested will produce response. Rather, it is currently designed to take in a list of compounds for which molecular MoR data exists and to rank-order them based on an individual patient's molecular concordance with the corresponding MR network hub. This is significant to note because the Bayesian assessment of predictive power likely underestimates the results as a consequence of this framework. Due to the limitations of the data (retrospective and no option for follow-up interventions), patients who are not given the best predicted treatment yet respond are nonetheless counted against the algorithm's performance - the algorithm is designed to pick the "best," but this does not preclude that "non-best" can work. Furthermore, this study is not meant to provide a ready-made assay format for the clinic. The goal of the study is to demonstrate that the field is now technologically capable of making network-based predictions of patient response using pre-treatment biopsies, and to encourage interest in developing the assays that would facilitate such assays for clinical applications.

These studies were also not designed to distinguish or characterize patients with long-standing AA, or *alopecia totalis* or *universalis*, which have been associated with poorer response rates to therapeutics in general, nor can we directly account for the relatively rare rate of spontaneous remission, since no reliable method currently exists for detecting it before the presentation of clinical remission. Further validation in the form of additional clinical trials or studies directly designed to inform methods such as this will be required to accurately include these parameters into a predictive model will greatly benefit not only validating this methodology further based on the framework we've presented in this study.

Regardless, this study establishes the precedent that network-based inference methods can be used to model and supplement drug matching decisions. The ability to provide quantitative estimates of efficacy for a panel of compounds will have significant influence on the overall success in treating diseases, as well as mitigate significant financial, temporal, and quality of life burdens and loss associated with patients being prescribed treatments that had little chance of working. In this study alone, in modeling three compounds, we observed variable response rates in each individual clinical trial. All outcome rates could have been systematically improved through the application of methods such as this beforehand.

While drugs with lower response rates – those with more specific and unique MoRs – will clearly benefit the most from such applications, even drugs with a high baseline efficacy can be enhanced by supplementing with these approaches by screening out even the rare non-responders early. The difference between 70% response to abatacept (random assignment) and 96% chance (assigned by algorithm), while perhaps less impressive than the other results, is nonetheless a significant improvement. More importantly, the quick identification of the 20% of patients who have a 0% chance of response to a treatment, before beginning treatment, would no doubt be appreciated by those individual patients. While the "stakes" of early identification in this disease context could be argued as minimal for clinical applications (e.g. "Why not just give ruxo to all patients?"), this technology has far-reaching consequences for the viability of precision medicine-based individualized testing, and these methods represent an invaluable tool in the future as we continue to develop new treatments for AA and identify feasible repurposing targets for known drugs.

Finally, such methods establish a novel framework and approach for translational biomarker development and have direct application potential to the clinic, directly affecting patient care and health. These proof-of-concept studies leverage whole skin biopsies of the scalp and bulk RNA sequencing to generate the predictive models. While this could already be directly implemented (if costly and time consuming) for patients before being prescribed medication, it also clearly defines large biomarker panels that can be used to refine more practical, compressed panels for testing. Perhaps with enough study and resources, more non-invasive biosampling could be adapted from this work to truly promote rapid, practical individualized patient matching.

4. Author Contributions and Acknowledgements

JCC and AMC designed this study and wrote the manuscript. JCC designed the core predictive model, implemented it, and validated it. ZD provided JAK/STAT single-targeting data via mouse model *in vivo* experiments. Funding for this project was provided by 1P50AR070588-01 and the Locks of Love Foundation. Special thanks to Eunice Lee, Eddy Wang, and Rolando Perez-Lorenzo for providing editing to the manuscript, and Daniel Smith for humoring the numerous applied mathematics discussions.

Declaration of Competing Interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Appendix A. Supplementary data

Supplementary data to this article can be found online at https://doi.org/10.1016/j.csbj.2021.08.026.

References

[1] Kennedy Crispin M et al. Safety and efficacy of the JAK inhibitor tofacitinib citrate in patients with alopecia areata. JCI Insight 2016;1. e89776.

[2] Jabbari A et al. Treatment of an alopecia areata patient with tofacitinib results in regrowth of hair and changes in serum and skin biomarkers. Exp. Dermatol. 2016;25:642–3.
[3] Mackay-Wiggan J et al. Oral ruxolitinib induces hair regrowth in patients with moderate-to-severe alopecia areata. JCI Insight 2016;1. e89790.
[4] Jabbari A, Cerise JE, Chen JC, Mackay-Wiggan J, Duvic M, Price V, et al. Molecular signatures define alopecia areata subtypes and transcriptional biomarkers. EBioMedicine 2016;7:240–7.
[5] Olsen E, Hordinsky M, McDonald-Hull S, Price V, Roberts J, Shapiro J, et al. Alopecia areata investigational assessment guidelines. National Alopecia Areata Foundation. J Am Acad Dermatol 1999;40(2):242–6.
[6] Dobreva A, Paus R, Cogan NG. Analysing the dynamics of a model for alopecia areata as an autoimmune disorder of hair follicle cycling. Math Med Biol 2018;35:387–407.
[7] Dobreva A, Paus R, Cogan NG. Toward predicting the spatio-temporal dynamics of alopecia areata lesions using partial differential equation analysis. Bull Math Biol 2020;82:34.
[8] Gilhar A, Keren A, Paus R. JAK inhibitors and alopecia areata. Lancet 2019;393 (10169):318–9.
[9] Chen J, Cerise J, Jabbari A, Clynes R, Christiano A. Master regulators of infiltrate recruitment in autoimmune disease identified through network-based molecular deconvolution. Cell Syst 2015;1(5):326–37.
[10] Margolin AA, Nemenman I, Basso K, Wiggins C, Stolovitzky G, Favera RD, et al. ARACNE: an algorithm for the reconstruction of gene regulatory networks in a mammalian cellular context. *BMC Bioinf* 2006;7(S1). https://doi.org/10.1186/1471-2105-7-S1-S7.
[11] Dai Z, Chen JC, Chang Y, Christiano AM. Selective inhibition of JAK3 signaling is sufficient to reverse alopecia areata. JCI Insight 2021;6.
[12] Phillips HS, Kharbanda S, Chen R, Forrest WF, Soriano RH, Wu TD, et al. Molecular subclasses of high-grade glioma predict prognosis, delineate a pattern of disease progression, and resemble stages in neurogenesis. Cancer Cell 2006;9(3):157–73.
[13] Lefebvre C et al. A human B-cell interactome identifies MYB and FOXM1 as master regulators of proliferation in germinal centers. Mol Syst Biol 2010;6:377.
[14] Mani KM et al. A systems biology approach to prediction of oncogenes and molecular perturbation targets in B-cell lymphomas. Mol Syst Biol 2008;4:169.
[15] Woo JH, Shimoni Y, Yang WS, Subramaniam P, Iyer A, Nicoletti P, et al. Elucidating compound mechanism of action by network perturbation analysis. Cell 2015;162(2):441–51.
[16] Pierre F, Chua PC, O'Brien SE, Siddiqui-Jain A, Bourbon P, Haddach M, et al. Pre-clinical characterization of CX-4945, a potent and selective small molecule inhibitor of CK2 for the treatment of cancer. Mol Cell Biochem 2011;356(1-2):37–43.