4-3-2020

Applying Systems Pharmacology to the Treatment of Chronic Illness Using Novel Scoring and Translational Methods

Spencer Richman
scr6045@rit.edu

Applying Systems Pharmacology to the Treatment of Chronic Illness Using Novel Scoring and Translational Methods

Thesis Submitted in Partial Fulfillment of the Requirements for the Degree of Master of Science in Bioinformatics

Spencer Richman[*†]

Advisory Committee:
Gordon Broderick[†‡]
Matthew Morris[*†]
Gary Skuse[*†]

[*]Gosnell School of Life Sciences, Rochester Institute of Technology
[†]Center for Clinical Systems Biology, Rochester General Hospital
[‡]Department of Biomedical Engineering, Rochester Institute of Technology

Rochester Institute of Technology
College of Science
Thomas H. Gosnell School of Life Sciences
Master of Science in Bioinformatics
04/03/2020

Dedication

This work is dedicated to the millions that have suffered and will continue to suffer from chronic disease.

Abstract

Chronic illnesses are poorly understood diseases that are often highly resistant to treatment. The prevalence and severity of these illnesses necessitates new methods for treatment development that diverge from the paradigm of one drug, one illness. Instead, multidrug interventions that utilize repurposable, previously approved drugs could be far more advantageous. In order to support this, a novel scoring framework and accompanying set of tools, collectively termed DrugAble, have been developed. DrugAble scores proposed, model-based treatment target solutions by analyzing drug-target interaction data and addressing the network complexity of these solutions. A*ctionability* scores that summarize the likelihood of a proposed target set constituting a pharmacologically accessible path to remission are generated. Additionally, DrugAble proposes combinations of repurposable drugs that can potentially be used in tandem to achieve remission. Here, DrugAble is demonstrated on molecular target solutions supporting an escape from Myalgic Encephalomyelitis / Chronic Fatigue Syndrome, a debilitating illness that affects up to 2.5 million Americans alone. DrugAble effectively discriminates between theoretical target sets and those that are clinically actionable using available drugs while simultaneously accounting for drug-target interactions and off-target effects of these drugs. This framework constitutes the necessary first steps to designing more effective treatments for chronic illnesses, with the ultimate goal of reducing the failure rate of clinical trials and the financial burden on both drug developers and patients. Most importantly, it opens new and more immediately accessible paths toward achieving remission and full recovery for those suffering from chronic illnesses.

Introduction

Chronic diseases are an increasingly studied though still incompletely understood set of illnesses with continual impacts on society. While exact definitions vary, chronic disease differs from traditional disease in that it is characterized by a long-term presence of symptoms, often greater than three months[1]. These diseases are usually multi-symptomatic and lead to a variety of health complications in affected persons[1]. According to the National Health Council, greater than 75% of all healthcare costs are due to chronic disease, and as of 2007, this was a total of 1.3 trillion U.S. dollars annually. It is projected that by 2023, these costs will reach upwards of 4.3 trillion dollars[1,2]. The economic costs speak nothing of the toll that these diseases take on the lives of those affected. Chronic diseases are often debilitating, and greatly reduce the quality of life of those suffering from such diseases. An increasing amount of multidisciplinary research is being conducted on how best to prevent, treat, and cure chronic diseases.

Over the past two decades, understanding of chronic disease has moved beyond hypotheses that attribute disease to single causes and that present as homogeneous phenotypes. It is increasingly understood that many chronic diseases are in fact highly heterogeneous in both etiology and presentation despite umbrella medical definitions[3–6]. This new understanding of chronic disease has led to an increased need for systems-level approaches to analysis and treatment of these illnesses. The field of systems biology arose from a need to analyze biological questions from a scope that far exceeds traditional analysis approaches. Viewing disease from the perspective of many co-interacting molecular and enzymatic pathways is a more robust and often more realistic perspective on disease[7,8]. Furthermore, because many chronic diseases lack a single identifiable cause, they are difficult to analyze using traditional methods. Systems biology allows researchers to analyze chronic disease from a wide angle and to consider multiple possible causes

of disease. When combined with other modern approaches to rapid analysis and treatment, systems biology becomes a powerful clinical tool[9,10]. However, modeling the full complexity of biological systems is computationally and temporally expensive, and thus the resolution of these models is often limited. By reducing the complexity of biological models into qualitative-- or discrete-- values and decisions, the size of the model can be increased even under relatively limited computational capacity. Instead of models containing dozens or even hundreds of parameters, models can be generated with as few as three parameters (a ternary model), each representing a directional, specific interaction between two system components[11]. For example, the effects of the addition of a drug to a system can be modeled as having one of three effects on each component of the system-- no interaction, inhibition, or activation. Using such models means that, even in the absence of kinetic data and numerous measurements that would otherwise be required, the ultimate state of the system can be calculated. This has the inherent advantage of being able to design larger models, but it comes at the cost of resolution. In reality, biological models are continuous functions. Nevertheless, discrete representations of biological systems can lead to robust and valuable results.

Such discrete models have been applied to solving and understanding hypothesized "steady states" responsible for chronic disease[12,13]. These steady states represent stable, homeostatic "basins" in the topology of a model, which, regardless of the beneficence of this state, the biological system maintains[12]. Movement into a new steady state can be triggered by perturbations that propagate through the system. In other words, these perturbations can trigger a movement from a stable, healthy state to a stable disease state-- chronic disease[12,13]. Discrete methods for analyzing the existence of and movement to these steady states work particularly well because

they do not require any knowledge of a system's kinetics. It does not matter when the system reaches the steady state, it only matters why it reaches such a state[12].

The development of these discrete logic models has enabled computation of minimally invasive drug interventions for the treatment of chronic disease. A Minimal Intervention Set (MIS) is defined by the least number of perturbations needed in order to induce a transition between steady states in a model[12]. In a clinical reality, this represents the least number of therapies that will induce the required changes. In the context of this work, an MIS constitutes a set of molecular target-action pairs. Mathematically, there are often multiple ways in which the system may come to rest in a given steady state, and multiple steady states in the system. Through the use of Monte Carlo methodologies and constraint satisfaction solvers, the behavior of a given system can be determined from any starting state, and the likelihood of the system following a certain path can be determined. This is further described in Craddock, et al., 2018. These same methodologies are applied in the development of treatments. By simulating the immediate effect of a treatment-- upregulation or downregulation of a component-- the system is then allowed to evolve towards a new steady state[12,13]. By repeating these simulations many times, the likelihood of a given treatment succeeding can be calculated. However, the bulk of existing work in this has focused on the identification of treatment targets regardless of how effectively or how specifically they might be modulated using known pharmaceutical compounds. Many if not most of these idealized treatments may translate poorly if at all to the clinical space as evidenced by the high failure rate of phase II clinical trials.

General approach

The current ethos underlying treatment design focuses on single-drug interventions that target a specific molecular process. Unfortunately, such treatments are often ineffective when

treating chronic disease because of the multifaceted causes of these diseases. As a result, clinical trials for treatments of many chronic illnesses fail[14–16]. The exceptionally long and costly nature of clinical trials[17] means that there is significant motivation to avoid dedicating resources to difficult-to-treat illnesses. In recent years, a new focus on the repurposability of drugs has arisen. Thousands of FDA approved drugs, originally for the treatment of one or more relevant conditions, may also be valuable for treating newly identified and newly understood illnesses[18,19]. This approach could be a significant boon to both pharmaceutical companies and sufferers of chronic disease, as time-to-market may be greatly reduced by implementing repurposing strategies. Nevertheless, there remains a need to both reduce the failure rate of clinical trials and to design new, efficacious treatments to chronic illnesses.

To address the failure of single-drug interventions, a paradigm shift is required. Multidrug interventions, where drugs work synergistically with one another may be more effective than individual drugs alone. This paradigm has already begun to take hold in some areas of medicine, especially for the treatment of HIV infection and various cancers. A similar paradigm, known as polypharmacology, has attempted to replace the single drug, single target philosophy with a "single drug, multiple target" philosophy[20]. However, this area of pharmacology is not amenable to repurposing of drugs and still does not address all of the challenges that the use of multiple drugs attempts to overcome. In complex chronic diseases, multidrug treatments may prove to be exceptionally effective.

There are important considerations that must be made when designing multidrug treatment courses. The treatments must be as minimally invasive as possible and the number of drugs, their frequency of dosage, and their side-effects should be as few as possible in order to reduce strain on the patient. While dosing parameters are outside the scope of this project, there remains a

challenge to effective treatment design that computational methods can help resolve. Off-target drug interactions can greatly complicate a treatment course through unintended effects to the system that were not known during the model generation phase.

Methods to detect off-target interactions have been widely employed in the field of drug discovery and design. Often, a variety of *in silico* procedures are used to identify these interactions. Work by Yera *et al.*, 2014 demonstrated a novel methodology combining natural language processing of patient packet information and computational analysis of the 2D and 3D structures of drug molecules using adapted established methods[21]. This data fusion technique produces a probabilistic model for how likely two drugs are to share molecular targets. The methodology was validated using the Structural Pharmacology Database and the chEMBL database. In 2017, Chartier *et al.*, using a combination of Protein Data Bank, PISCES, and DrugBank data, performed binding-site similarity analyses for 400 drugs and 7,895 different proteins. Chartier identified multiple drug-repurposing candidates as well as a large set of interactions that may explain the side-effects of certain drugs[22]. These works successfully demonstrate valid means for the computational prediction of off-target drug interactions. While these applications have traditionally been applied to drug discovery and design, similar methods are justifiably applicable to treatment design as well. Because predictive methods based on physical interactions and structural models are computationally intensive and limited to proteins with solved structures, other approaches may be preferable. The application of drug-target information for multidrug treatment design has precedent. Torres et al, 2016, analyzed pairs of drugs in DrugBank and developed scoring methods for predicting synergistic effects between pairs of drugs while also accounting for off-target effects[23]. However, their work focused specifically on pairs of drugs and

on developing signed networks of drug-target effects without focus on scoring differing solutions, repurposing drugs, or predicting treatment combinations.

In order to address the above-mentioned translational gaps, the current work is focused on aligning model-predicted molecular and cellular targets with pharmaceutical agents already in clinical practice as well as known compounds biochemically catalogued but not yet approved for use. To meet this goal, quantitative criteria have been developed to assess the relative ranking of treatment target sets on the basis of their off-target effects, the interchangeability of compounds, expected downstream effects and the consensus among models in the robustness of predicted response. This theory has been implemented in the form of a computational framework known as DrugAble, which mines existing knowledge-bases for drug-target interaction data, links these data to targets in a provided model, ranks the solutions to the model, and translates high-scoring solutions into pharmaceutical interventions using repurposable compounds.

Methods

Drug Identification

Effective implementation of the proposed scoring methodology requires maximal data on drug-target interactions. There are many web databases with drug-target information, including DrugBank, PharmGKB, and others. However, because effective data fusion from these sources, which do not utilize a standard format nor standard naming conventions, is difficult an alternative was sought. Instead, an API interface to the Elsevier Pathway Studio Service, dubbed DTQuery, is implemented. The Pathway Studio database contains a massive amount of drug-target information collected from literature mining of 3.5 million full texts and 24 million PubMed abstract, as well as putative interactions imported from Reaxys Medicinal Chemistry. DTQuery accepts input as a JSON file describing the modulations required for each target in all possible

MIS candidates and a file containing target names, their alternative names, and identifiers. Using this information, a series of queries is performed to the web database. Query 1, provided in a pseudo-SQL form, searches for all small-molecule drugs that modulate any of the provided targets. Query 2 then iterates over the results of the first query, returning all known targets of each identified small-molecule drug. This provides the necessary information on off-target effects. Every drug-target interaction that is identified is written to a local flat file.

Query 1: Drugs that Modulate Each Target

```
SELECT RELATION
WHERE UPSTREAM_NEIGHBOR
(SELECT NAME = target
AND TYPE = ()
AND DOWNSTREAM_NEIGHBOR
(SELECT TYPE = 'Small Molecule'))
AND TYPE = (Regulation, DirectRegulation, Binding, Transport))
AND EFFECT = (POSITIVE, NEGATIVE)
AND ORGANISM = 'Homo sapiens')
```

Query 2: Find all Off-Target Effects

```
SELECT RELATION
WHERE DOWNSTREAM_NEIGHBOR
(SELECT NAME = DRUG)
AND UPSTREAM_NEIGHBOR
(SELECT TYPE = (Regulation, DirectRegulation, Binding,Transport)
AND EFFECT = (POSITIVE, NEGATIVE)
AND ORGANISM = 'Homo sapiens')
```

Candidate Solution Scoring

The scoring of input solutions is a critical component of the DrugAble pipeline. A novel scoring function, henceforth known as Actionability (α), is designed to be a quantitative measure of the beneficence of a solution. Each of the five component scores summarize key information about both the solution and available drugs that may satisfy the parameters of the solution. Additionally,

two tuning parameters are defined that allow for weighted adjustment of the contribution of *cardinality* and the *antagonist ratio*, described below. The composite nature of Actionability allows fine-tuning to account for both the complexity of the solutions and the nature of the drugs used to achieve the solutions.

In the following equations, a model is defined as a set of logical vectors each representing a solution to the discrete logic model being analyzed. Iverson bracket notation is used for logical operations, where conditions within square brackets evaluate to 1 if true and 0 if false. Finally, set cardinality is defined by the vertical line enclosure (||).

The first component of α is *cardinality* (*C)*. Cardinality is defined in this instance as the sum of non-zero elements of the polarity vector that defines a candidate solution to a model. The formal definition for the cardinality of a vector *s* is provided below in Equation (1) :

Equation 1:

$$C(s) = \sum_{i=1}^{n} [s_i \neq 0]$$

where n is the number of targets in the model, s_i is the polarity of the target at index i, and 0 is a neutral polarity. It is generally best to promote candidate solutions with minimal cardinality. Candidate solutions with greater cardinalities will likely require a greater number of drugs in order to achieve a remissive state, which introduces a greater risk of destabilizing off-target effects and negative side effects for patients.

The second component score, the *antagonist ratio* (*V*) is defined as the ratio of antagonist actions to agonist actions in a candidate solution. It is generally accepted in pharmacology that it is easier to downregulate a drug target than to upregulate one []. The antagonist ratio has been incorporated in order to promote candidate solutions that require less upregulation of targets, and

thus are more likely to have available drugs. The formal definition for a candidate solution *s* is provided below in Equation (2):

Equation 2:

$$V(s) = \frac{p}{q} \begin{cases} p = \sum_{i=1}^{n} [s_i = 1] \\ q = \sum_{i=1}^{n} [s_i = 2] \end{cases}$$

where *n* is the number of targets to be modulated, s_i is the polarity of the target at index *i*, *p* is the sum of antagonist actions, *q* is the sum of agonist actions, 1 arbitrarily represents an antagonist action and 2 arbitrarily represents an agonist action. Maximizing the antagonist ratio is ideal.

The third component score, *repurposability* (*R*), is a measure of the likelihood that a candidate solution can be translated into a true pharmaceutical intervention. It is a function of the number of drugs that exist for each target. The greater the number of available drugs, the more likely that a valid drug combination can be generated. *Repurposability* also serves as a filtering score. Any candidate solution with no available drugs automatically receives an *R* score of zero, which by design floors α to zero. The formal definition for the *repurposability* of a candidate solution *s* is provided in Equation (3):

Equation 3:

$$R(s) = \left(\prod_{i=1}^{n} |\hat{s}_i| \right)^{1/n}$$

where *n* is the number of targets to be modulated and $\hat{s}_i$ is the set of all drugs that act with the correct polarity on s_i. Solutions with higher *R* are likely to have more varied combinations of drugs available when generating interventions when compared with lower *R* values.

The fourth component of *α*, the *target score* (*T*) effectively measures the mean number of targets a given intervention developed from a candidate solution is expected to have. This measure is agnostic to both whether interactions are on or off-target and to the polarity of the interactions. The formal definition for the *target score* of a candidate solution *s* is provided in Equation (4):

Equation 4:

$$T(s) = \left(\prod_{i=1}^{n} min(K_i) \right)^{1/n}$$

where *n* is the number of targets to be modulated and K_i is the set of sums of interactions for every drug that also interacts with target *i*. Thus, *T(s)* is the geometric mean of the minimum possible number of target interactions across all targets in *s*. Solutions with a lower *T* are more likely to yield drug interventions with fewer off-target effects, and so it follows that minimizing *T* is ideal.

α incorporates the four component scores as well as two tuning parameters, τ and γ, which allow for adjusting the weights of *V* and *C* respectively. Because *V* and *C* are derived from the structure of the MIS in contrast to the data-driven *R* and *T* scores, the ability to adjust their contributions to the final score can be useful. Higher values of τ and γ place scoring emphasis on

ideal MIS structure. Being able to adjust τ and γ independently from one another allows for even finer rank-tuning. The formal definition of α for a given candidate solution s is provided in Equation (5):

Equation 5:

$$\alpha(s) = \frac{R(s) * V(s)^{\tau}}{C(s)^{\gamma} * T(s)}$$

where *R*, *V*, *C*, and *T* are the same as defined in Equations (1:5).

While individual component scores can be useful when gauging the viability of a candidate solution, a single summary measure is far wieldier. In addition, the ability to weight α with the provided tuning parameters means that implementations of the scoring system can be fine-tuned for what is considered most desirable in each analysis. In any circumstance, a more maximal α for a solution is considered best. Final actionability scores are normalized between 0 and 1 for ease of interpretation. It should be noted that neither α nor any of its component scores can be directly compared with the same scores from a non-structurally identical model, nor across differing drug-target datasets. These scores are only valid in the context of the specific models and datasets used to generate the scores. There is opportunity for future research into developing a normalized score that can compare scores across different datasets and models.

Intervention Generation

The final major component of the DrugAble pipeline translates high-scoring candidate solutions into proposed clinical interventions. These proposed interventions must satisfy two restrictions. First, they must act upon, with correct polarity, every target prescribed to be modulated by the candidate solution. Second, they must use no more drugs than the cardinality of the candidate solution. Interventions that meet these requirements are considered valid.

Exhaustive generation of interventions is best-formulated as an NP-hard constraint-satisfaction problem[24]. In order to overcome the challenges brought by this, a heuristic approach using search-space reduction and ranking has been devised. Search-space reduction to identify viable drugs is performed based on three requirements. First, the drug must modulate at least one prescribed target in the candidate solution with correct polarity. Second, the drug must not modulate any other target in the solution with incorrect polarity. Third, any drug whose mean number of off-target effects is greater than one standard deviation from the mean number of off-target effects for all drugs is excluded. Ranking is similarly performed in two steps. In the first ranking operation, drugs with a higher number of on-target effects are promoted. In the second ranking operation, drugs with more off-target effects are demoted. If at this point there are any drugs whose number on-target effects is equal to the cardinality of the candidate solution, a rare single-drug intervention has been identified. Most often, this is not the case.

Multi-drug interventions are then generated using an iterative combinatorial algorithm. First, all combinations of size 2:*m* from drugs *n*, formulated as *n Choose 2*:*m*, are generated, where *m* is no larger than the cardinality of the candidate solution. From this, a large random sample of combinations are selected. Second, all combinations whose total number of on-target effects are less than the cardinality of the candidate solution are discarded. Third, combinations containing redundant matches, such as two or more drugs modulating all of the same targets, are culled. The resulting set of valid drug combinations is then output. Use of a heuristic “filter-rank-filter” methodology allows for the generation of valid interventions in far less compute-time than a constraint-satisfaction solver or exhaustive guess-and-check strategies.

Modeling Myalgic Encephalomyelitis / Chronic Fatigue Syndrome

Myalgic Encephalomyelitis / Chronic Fatigue Syndrome is a severe chronic illness that affects millions of individuals across the world. Over a million in the United States alone suffer from illnesses that can be categorized as ME / CFS . ME / CFS, while highly complex, is generally categorized by long periods of severe bodily fatigue rendering sufferers unable to function effectively. In addition to fatigue, sufferers often experience debilitating whole-body pains and mental deficits. ME / CFS, affects upwards of one-percent of the world's population, of which the overwhelming majority are women[25]. Both the etiology and presentation of ME / CFS are poorly understood, and in fact may comprise a set of similar but distinct conditions. Nevertheless, ME / CFS is increasingly understood in part to be the result of immune dysregulation, likely induced by viral infection[26,27].

Morris *et al.*, 2019 developed a regulatory model of ME / CFS based on 28 molecular targets, 214 regulatory interactions, and 17 immune markers measured from a cohort of 88 females with diagnosed ME / CFS[28] . These targets are outlined in Table 1. These 28 targets form a highly interconnected regulatory network between the immune and endocrine systems, as well as the HPA and HPG axes. Sex hormones exercise significant control over the immune system, especially in females[29], and this multi-axis network is believed to become dysregulated following severe viral infection, such as by the Epstein-Barr virus[30].

Target	Class	Function
Norepinephrine	Hormone	Stress response modulation
Estrogen	Hormone	Development and regulation of female reproductive system
Dopamine	Hormone	Motor control, hormone control
Cortisol	Hormone	Blood sugar modulation, immune system suppression
Luteinizing Hormone	Hormone	Testosterone production (males), ovulation (females)
Follicle-stimulating hormone	Hormone	Body growth and development, pubertal maturation, reproductive modulation
Gonadotropin releasing hormone 1	Hormone	FSH modulation, LH modulation
Corticotropin releasing hormone	Hormone	ACTH synthesis modulation
Progesterone	Hormone	Menstrual cycle, pregnancy, embryogenesis (females), testosterone production (males)
Stress	Physiological response	Survival instinct promotion
Natural Killer Cell	Cell	Response to viral infection
B-Cell	Cell	Antibody & cytokine secretion, antigen recognition
Interferon Gamma	Cytokine	Macrophage activation, MHC induction, response to viral infection
Tumor Necrosis Factor	Cytokine	Cell proliferation and differentiation, cell death
IL1A	Cytokine	Inflammatory response, T_h2 cell activation
IL1B	Cytokine	Inflammatory response
IL2	Cytokine	Immune response modulation

IL4	Cytokine	T_h2 cell induction
IL5	Cytokine	B-cell growth stimulation
IL6	Cytokine	Multifunctional, pro/anti-inflammatory response
C-X-C motif chemokine ligand 8 (CXCL8)	Chemokine	pro-inflammatory response, neutrophil recruitment
IL10	Cytokine	B-cell proliferation, T_h1 inhibition
IL12	Cytokine	T_h1 cell induction
IL13	Cytokine	Allergic immune response
IL15	Cytokine	Anti-apoptotic response
IL17A	Cytokine	Proinflammatory response, NFKB & MAPK regulation
IL23	Cytokine	Proinflammatory response, T_h17 mediator

Table 1: Parameters of the Myalgic Encephalomyelitis / Chronic Fatigue Syndrome Model
Twenty-Eight hormones, cytokines, chemokines, and immune cells form the interaction network model of ME / CFS. All drugs identified by DTQuery directly act upon one or more of these targets. However, no drugs were identified that directly target NK-Cell or stress.

Currently, no effective interventions, pharmaceutical or otherwise, exist for the treatment of ME / CFS[31]. In order to address this issue, MIS candidate solutions from five selected models of ME / CFS were pooled. These MIS candidates serve here as data for the proof-of-concept implementation of DrugAble. Preprocessed solutions were fed to DTQuery in order to identify relevant drug-target interactions, and this dataset of interactions was used to perform actionability scoring and intervention generation for this solution set.

Results

In total, 44,845 drug-target interactions (19,013 agonist, 16,367 antagonist) were identified using target-action pairs described by 2,515 MIS candidates. This includes traditional small-molecules drugs, lead compounds, monoclonal antibodies, and query contaminants. Query

contaminants include elemental dimers, nonspecific hits such as DMSO, and other non-pharmaceutical compounds that were identified during querying. The effect of contaminants on dataset quality is demonstrated in Figure 1A, where the Ca^{2+} ion has 632 annotated interactions, but is not an actual pharmaceutical. In figure 1B, ambiguous interactions were removed from the dataset, preserving only interactions annotated as "positive" or "negative." This filtered Ca^{2+} and other contaminants from the dataset, tightening the distribution of counts of targets per drug significantly. In Figure 1C, the distribution of drugs can be seen per in-model target. Candidate solutions where targets with high opportunity for repurposability (*R* score) are more likely to receive a higher actionability score, though this is also dependent on the other component scores of Actionability. In this dataset, TNF, IL6, CXCL8, and IL1B appear to be exceptionally druggable. As expected, most drugs identified for these targets act as inhibitors. Oppositely, neither NK-cell nor stress are druggable. NK-cells, because of their central and essential role in the innate immune system, would almost never be intentionally targeted by a pharmaceutical. Stress, on the other hand, is difficult to target when it is treated as a single condition. Stress is itself multifaceted and not a single molecular target. It would therefore be rare for a drug to be annotated as directly targeting stress. Figure 1C also facets the total number of drugs identified per target by whether those drugs inhibit or promote the target. Across all targets, there are 1.8x more drugs that inhibit these targets than there are that promote it, justifying the inclusion of the antagonist ratio as a component score of actionability.

The propensities for each target in the ME / CFS model to be inhibited or promoted can be seen in Figure 2. Interestingly, NK-cell and stress, the two targets for which no drugs were identified, are never modulated in any of the candidate solutions. ACTH, IL2, IL4, and IL5 are also never modulated. No single target is modulated in every candidate solution, suggesting the

presence of multiple subgraphs that lead to the same outcome. In Figure 3, the number of solutions by cardinality is divided per model. While model 3 produced the majority of total solutions (2005), it produced almost exclusively high-cardinality solutions. Combined, 97 percent of solutions had a cardinality of five or greater. Only 2.3 percent of solutions were of cardinality four or less, with none below a cardinality of three. The lack of low-cardinality solutions reinforces the highly complex nature of ME / CFS, and demonstrates the necessity of including cardinality as a component score of actionability.

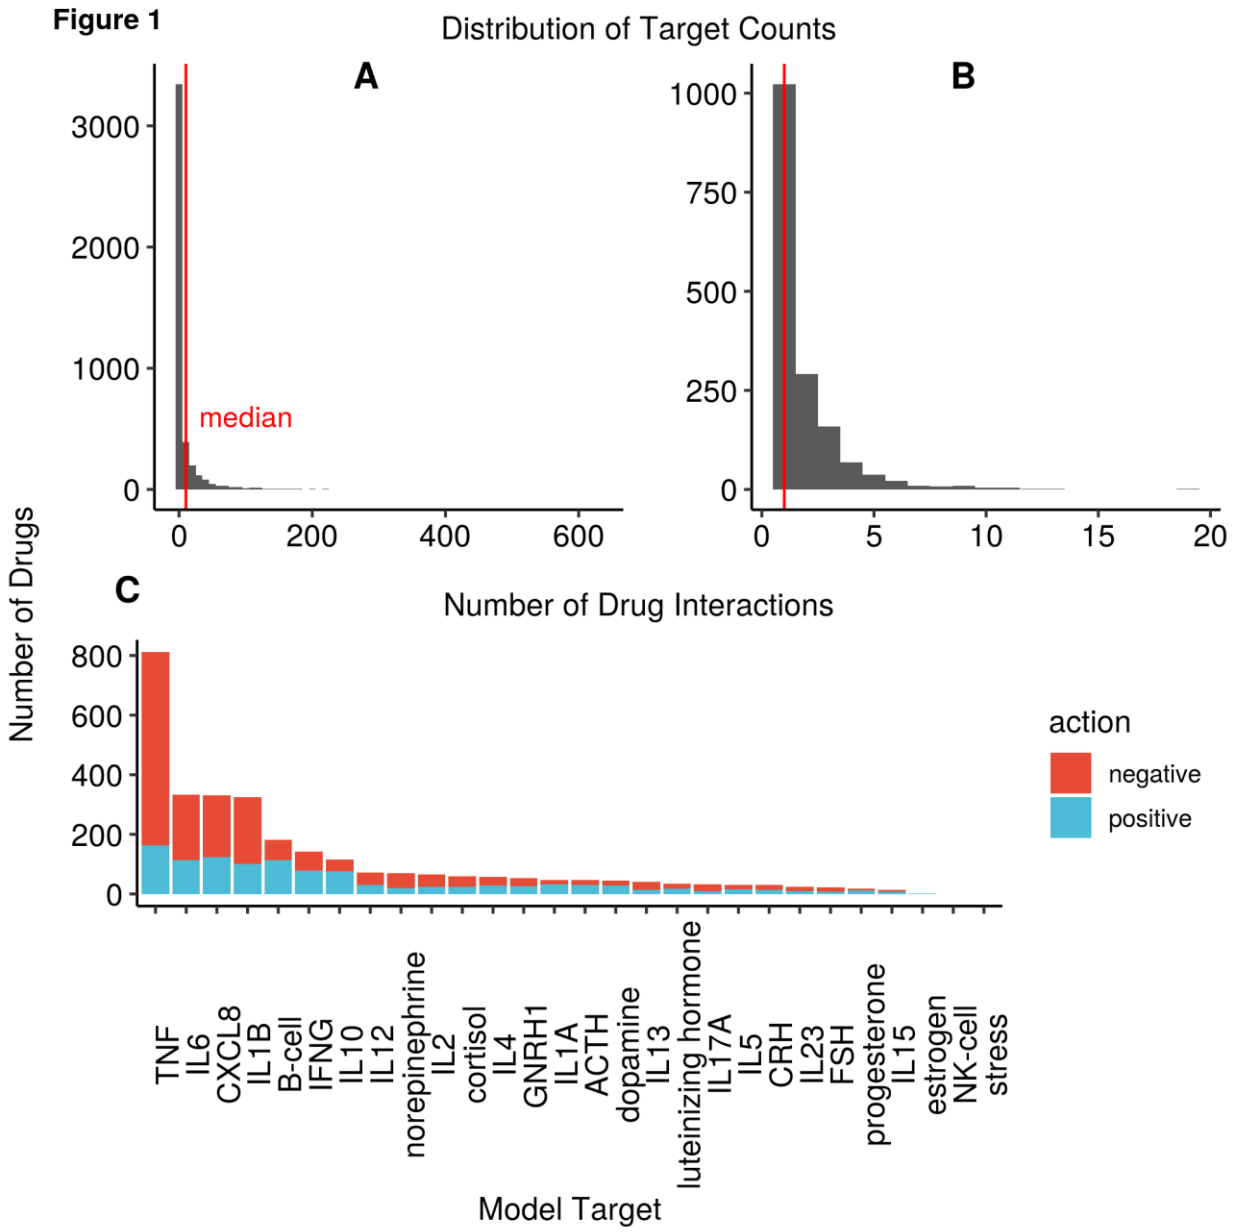

Figure 1: Drugs Identified by DTQuery
A summary of the drugs and compounds identified by querying the Elsevier Biological Knowledgebase. In (A), the distribution of all 44,845 relations is presented, including non-specific annotated interactions such as interactions labeled as "unknown." The median number of interactions is 1.0 interaction / drug. In (B) ambiguous interactions have been removed, leaving only interactions annotated as "positive" or "negative." Additionally, interactions with targets outside of the model have been removed. In (C) a breakdown of the polarity of interactions by target is provided. Note that *NK-cell* and *stress* are truly zero.

Figure 2

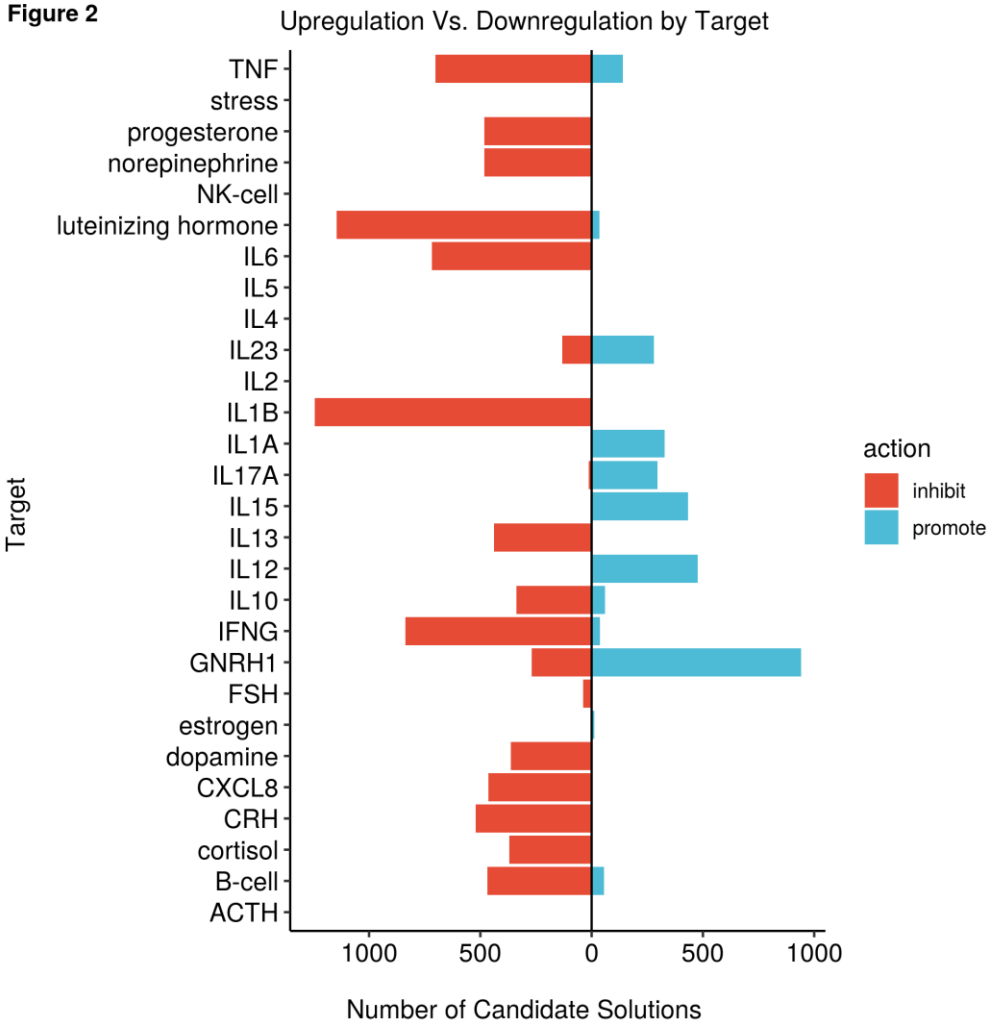

Figure 2: Upregulation and Downregulation of Targets
The majority of targets are inhibited more often than they are promoted, with the notable exceptions of GNRH1, IL1A, IL12, IL15, and IL17A. IL2, IL4, IL5, NK-cell, stress, and ACTH are never modulated.

Figure 3

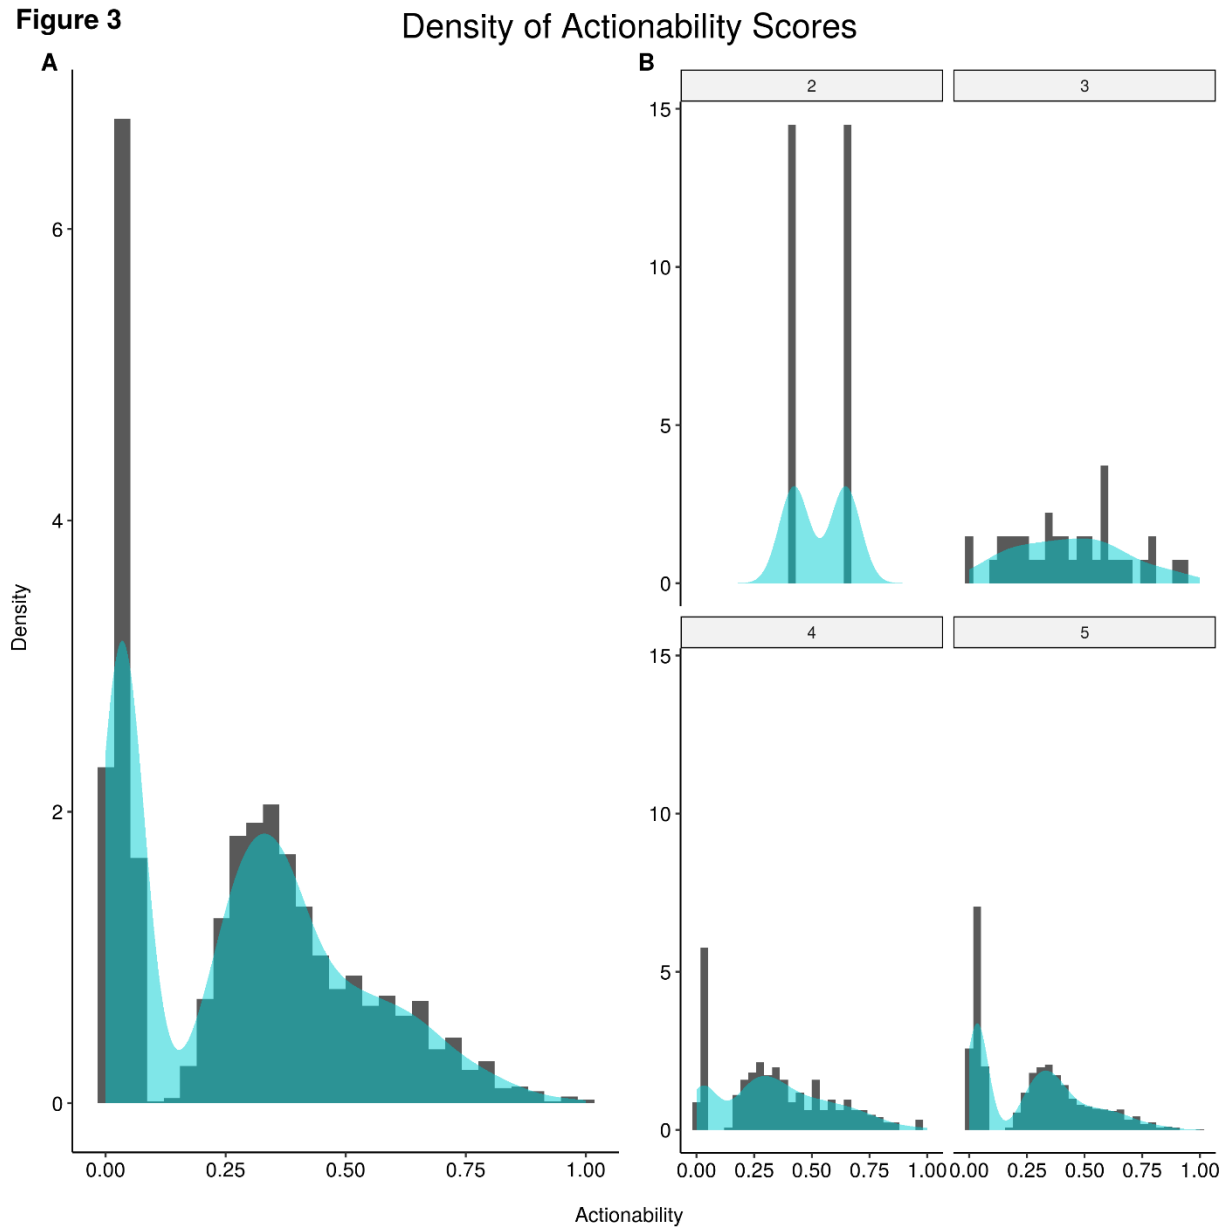

Figure 3: Density of Actionability Scores
A) The distribution of all actionability scores is presented. **B)** The actionability scores are separated by cardinality, from 2 to 5. Actionability is normalized from 0.0 to 1.00, where 1.00 is the most actionable solution. The blue highlighted area is a best-fit curve to describe the distribution of scores.

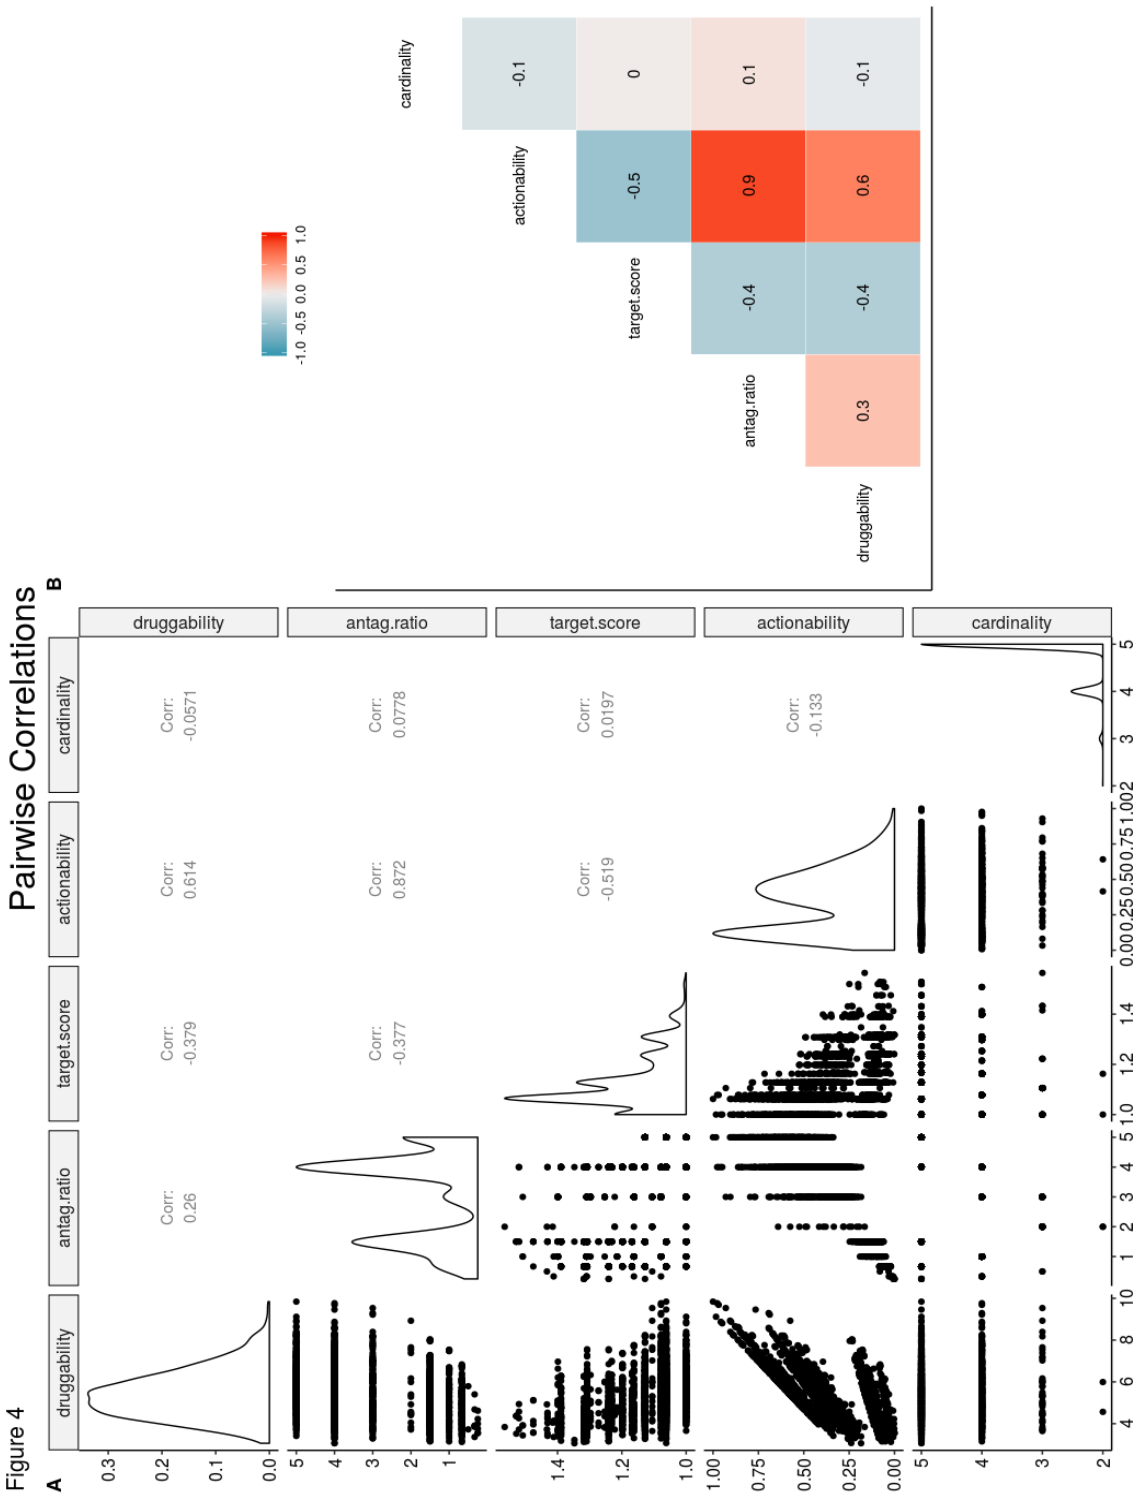

Figure 4: Pairwise Correlations of Component Scores

A) Pairwise scatterplots for each component score of actionability verses each other component, as well as the total actionability score. Diagonals are the distribution of each score, and the upper triangle shows the Pearson correlation for the corresponding cells. **B)** The same correlations shown in the upper triangle of A, but color-coded by magnitude.

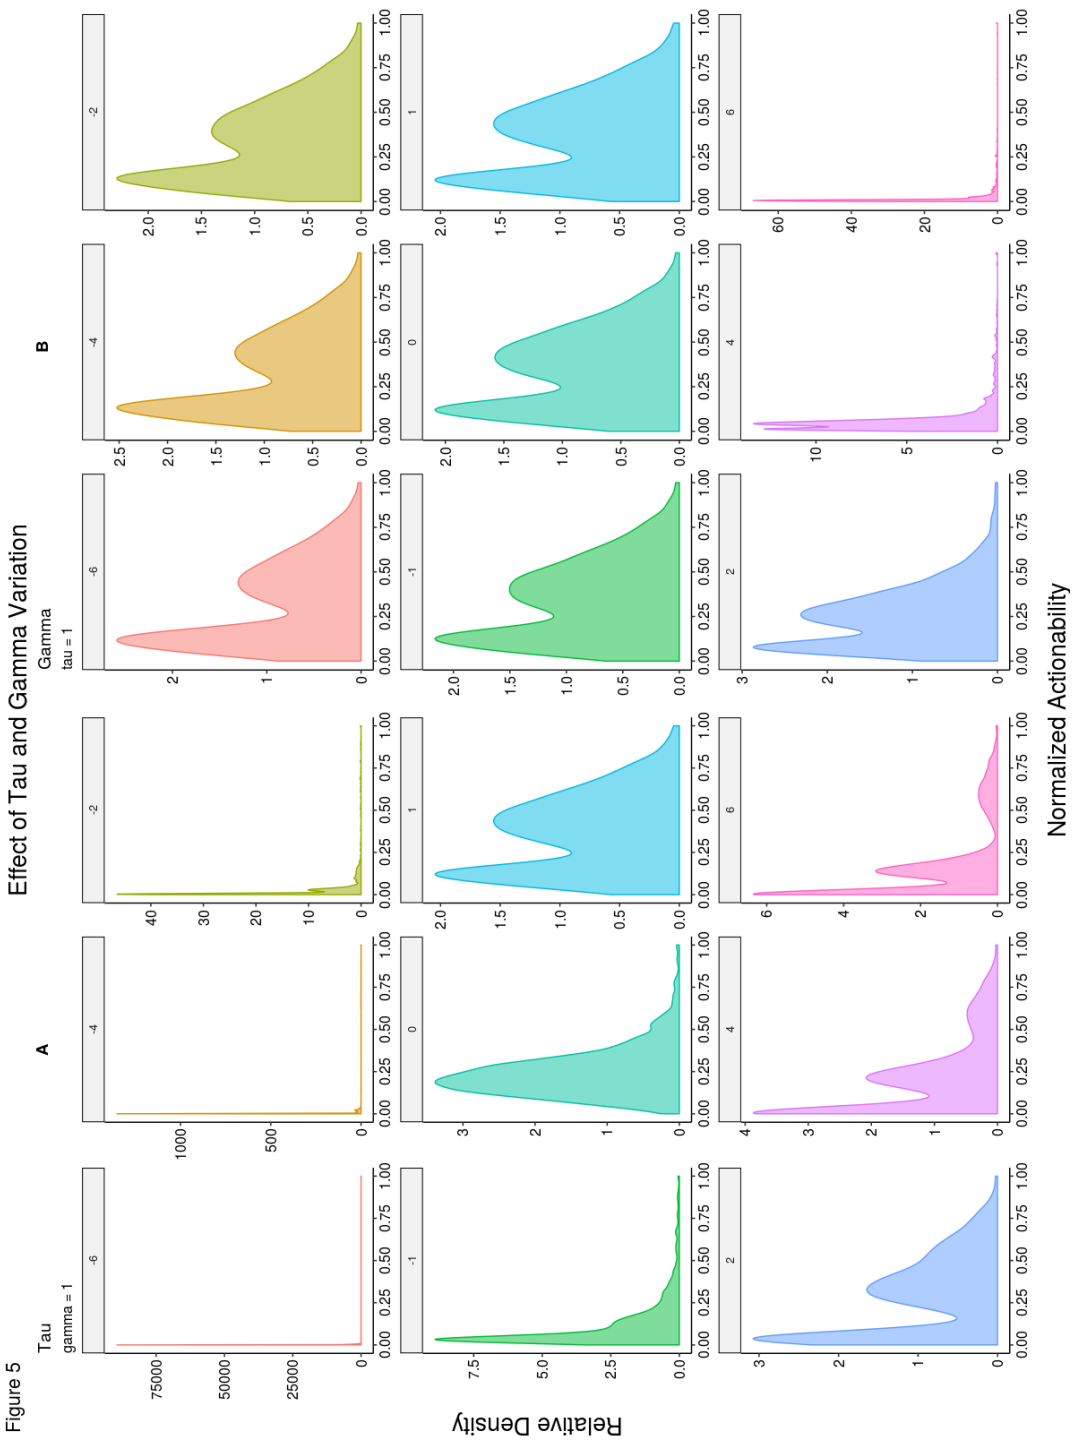

Figure 5: The Effects of Tau and Gamma Tuning Parameters on Actionability
A) Negative values of tau invert the antagonist ratio metric, instead prioritizing solutions with high agonist ratios. Positive values increase the weight of antagonist ratios on actionability. **B)** Negative values of gamma should not be used, as they prioritize high-cardinality solutions. Positive values increase the weight of cardinality on actionability. For both tau and gamma, values between 1 and 4 are recommended.

Actionability is not a statistical measure. It does not attempt to assign a probability of success or failure to a solution, only to rank solutions by their ability to be translated into pharmaceutical interventions in comparison to all other solutions available, using known drug-target interactions and knowledge about the structure of the solution. A normalized actionability score of 0.9 for a given candidate solution does not mean that this solution has a 90 percent chance of successfully being translated or succeeding as a pharmaceutical intervention. Rather, an α of 0.9 says that a solution is 90 percent more likely to be a good solution, as defined by the standards of an ideal solution, than the lowest-ranking solution.

Nevertheless, it is important to verify that actionability is a meaningful measure of the ideality of a solution. Before the implementation of actionability, the only measure for assessing the goodness of a candidate solution was cardinality. Cardinality, while an important metric-- hence its inclusion in actionability-- does not provide enough information alone to accurately score MIS candidates. As figure 4 demonstrates, no component score alone is able to provide the same level of discrimination as actionability is in total. While the antagonist ratio has a strongly positive correlation (0.9) with actionability, this correlation is somewhat spurious. Despite a clear upward trend in actionability with increasing antagonist ratio, which is expected as per the function of the antagonist ratio in calculating actionability, the possible values of actionability vary greatly at each level of antagonist ratio. The highest antagonist ratio predicts normalized actionabilities ranging from 0.215 to 1.0, a very large range. The distribution of each score is also worth noting. Actionability is exceptionally bimodal, with a very high concentration of very low-actionability scores. The bimodal nature of this distribution is likely caused by the filtering effects of the component scores, which rather harshly cause solutions with even one or two non-ideal component scores to drop out. The antagonist ratio distribution is also bimodal, with values concentrated at

either end. This suggests that it is common for solutions to have either many antagonized targets or many agonized targets, and far less often a comparable number of both, but it is important to note that this can vary greatly between models and illnesses. The moderate negative correlation (-0.4) between repurposability and target score is also to be expected, because with an increasing number of available drugs (I.e., a higher repurposability), the target score is likely to increase due to an increasing number of off-target effects from those drugs.

The strictness of actionability can be adjusted with the tuning parameters τ and γ. These two parameters act as exponential scaling factors for the antagonist ratio and cardinality respectively. The effect of varying levels of τ and γ on the final actionability score can be seen in Figure 5. In Figure 5A, the distribution of actionability scores at τ values between -6 and +6 are shown. Any value of τ less than one increases the strictness of actionability greatly. The bimodal distribution of actionability seen in positive values of τ disappears. For this dataset, any value of τ below -2 is largely uninformative because most scores become effectively zero. Increasingly positive values of τ also impart strictness on actionability, though this effect is less extreme than negative values. As τ increases from 1, the bimodal distribution in this dataset is preserved, but is shifted into a narrower range toward the lower end of actionability scores. The behavior of γ (Figure 5B) is distinctly different from that of τ. Negative γ values have a much less drastic effect on the distribution of actionability scores, though increasingly negative values do lead to a slightly greater density of mid-level scores. Increasingly positive γ values have a similar but less-pronounced effect on the distribution that increasingly negative values of τ do. In practice, in datasets with a high number of small-molecule drugs, a higher τ will cause greater separation in actionability between candidate solutions with high antagonist ratio and those with low ratios. This may be preferable because small-molecule drugs tend to be antagonists. In the same scenario, a

high γ will penalize candidate solutions with a high likelihood of generating interventions that have many off-target interactions. This may be preferable because it will indirectly penalize the use of promiscuous small-molecule drugs. Conversely, if a dataset is rich in agonist drugs, or a model is sparse in candidate solutions with antagonist actions, reducing or even inverting the weighting of these scores can be beneficial. Combined, these two parameters can be used to tune actionability so that it remains flexible under different assumptions of the ideality of candidate solutions.

The generation of proposed pharmaceutical interventions is the final major function of DrugAble. It is also the most computationally demanding. In order to avoid the complexity and pitfalls of implementing a constraint satisfaction solver to identify combinations, a filter-rank-filter methodology, described above in the methods section, was implemented. The top ten candidate solutions were selected for translation to proposed interventions based on their actionability scores. For each of the ten candidate solutions, all possible interventions of two and three drugs were generated. Because the generation of combinations is performed iteratively down the supplied list of drugs, the resulting combinations are implicitly sorted by the least number of off-target effects and highest number of on-target effects. For each of the top ten solutions ranked by actionability, the top 20,000 combinations were checked for validity. Table 2 shows the target-action pairs for the ten highest-ranking solutions, and Table 3 shows the breakdown of valid combinations for those. Tables 4 and 5 are offered in contrast, showing the same data but for the 10 lowest-ranking solutions. Table 5 shows the top-ranking interventions for each of the 10 highest-ranking solutions.

Actionability	IL1B	IL16	CXCL8	IL10	TNF	B-Cell	Luteinizing Hormone
1.00	1	1	1	1	1	0	0
0.98	1	1	1	0	1	0	1
0.98	1	1	0	1	1	0	0
0.97	0	1	1	1	1	0	0
0.96	1	1	0	0	1	0	1
0.95	0	1	1	0	1	0	1
0.93	0	1	0	1	1	0	0
0.91	1	1	0	1	1	1	0
0.90	0	1	0	0	1	0	1
0.90	0	1	1	1	1	1	0

Table 2: Best Solutions by Actionability
Each of the best ten solutions, ranked by actionability. A *1* represents an antagonistic action, a *0* represents no action, and a *2* represents an agonistic action (none present).

Actionability	valid	M=2	M=3	Mean OT
1.00	483	57	426	69.6
0.98	88	6	82	62.4
0.98	2254	173	2081	10.1
0.97	2363	171	2191	12.4
0.96	131	27	104	15.6
0.95	135	27	108	17.5
0.93	3000	1310	1690	6.9
0.90	158	21	137	6.9
0.90	105	1	104	21.0

Table 3: Interventions Generated for Top Solutions
A breakdown of valid interventions generated for each of the top ten solutions, ranked by actionability. M=2 refers to interventions of size 2, and M=3 refers to interventions of size 3. Mean OT is the mean number of off-target effects across all generated interventions.

Actionability	IL1B	IL15	IL1A	IL12	IL17A	IFNG	B-Cell	Cortisol	GNRH1	Luteinizing Hormone
0.010	0	0	2	0	2	0	0	1	2	0
0.010	0	0	0	2	2	0	0	1	2	0
0.010	0	2	0	0	2	0	2	0	1	0
0.008	1	2	2	2	0	0	0	0	2	0
0.007	0	2	0	0	2	0	0	1	2	0
0.003	0	0	2	2	2	0	0	1	2	0
0.003	0	0	2	2	2	1	0	0	2	0
0.002	0	2	0	0	2	0	2	0	1	2
0.000	0	2	2	0	2	0	0	1	2	0
0.000	0	2	0	2	2	0	0	1	2	0

Table 4: Worst Solutions by Actionability
Each of the worst ten solutions, ranked by actionability. A *1* represents an antagonistic action, a *2* an agonist action, and a *0* represents no action.

Actionability	Drug 1	Drug 1 Targets	Drug 2	Drug 2 Targets	Off-Targets
Solution 1					
1.00	Rosmarinic acid	IL1B,I L16,TNF, CXCL8	IC14	Il6, IL10	12
...	...	...	Oxprenolol	IL10	14
...	...	...	norcantharidin	IL10, TNf	33
...	...	...	eplerenone	CXCL8, IL10	21
Solution 2					
0.98	Rosmarinic acid	IL1B,IL16,TNF, CXCL8	Levonorgestrel	Luteinizing Hormone	22
...	...	...	Substance P Antibody	Luteinizing Hormone	10
...	...	...	Galantide	Luteinizing Hormone	12
Solution 3					
0.98	Lipid A Derivatives	IL1B, IL6, TNF	IL12 Antibody	IL10	1
...	...	...	Taurolidine	IL10	1
...	Exosurf Neonatal	IL1B, IL6, TNF	IL12 Antibody	IL10	1
...	...	....	Taurolidine	IL10	1
Solution 4					
0.97	Pulegone	IL16, CXCL8, TNF	IL12 Antibody	IL10	1
...	...	...	Taurolidine	IL10	1
...	Ammonium Acetate	IL16, CXCL8, TNF	IL12 Antibody	IL10	1
...	...	...	Taurolidine	IL10	1
Solution 5					
0.96	Lipid A Derivatives	IL1B, CXCL8, TNF	Substance P Antibody	Luteinizing Hormone	0
...	Exosurf Neonatal	IL1B, CXCL8, TNF	...	...	0

Table 5: Proposed Interventions
A subset of the generated, predicted interventions for the first five solutions. Ellipses indicate that the cell is unchanged from the cell directly above.

Discussion

DrugAble implements flexible quantitative measures that allow for the ranking of idealized candidate solutions to chronic illnesses. The actionability summary measure accounts for both known drug-target interactions and the complexity of candidate solutions. Importantly, actionability can be tuned to weigh its component scores differently. This means that users can define, to a degree, what defines an actionable solution. This is important because no two illnesses, nor the drug datasets used to examine these illnesses, are the same. While actionability by default discourages solutions with high numbers of agonistic requirements, it is not infeasible that a researcher may have a drug-target dataset rich in agonistic drugs. In this case, lowering or even reversing this weighting may be desirable. For ME / CFS targets and the accompanying drug dataset mined by DTQuery, agonistic actions were often unsatisfiable or undesirable. This is evidenced by the total lack of agonist actions in the top ten highest-scoring solutions, and the heavy presence of agonist actions in the lowest scoring solutions, seen in Tables 2 and 4. While the actionability scoring was tuned to intentionally lower the rank of agonist actions, this was clearly justified based on the inability to develop any valid interventions for the lowest-ranking solutions. Most importantly, actionability is a marked improvement over cardinality alone, the previous metric used to judge the viability of a solution. As can be seen in Figure 3B, low-cardinality solutions do not necessarily represent highly actionable ones. The range of actionability scores for each value of cardinality is often large. If solutions are ranked on cardinality alone, many poorly clinically actionable solutions may rise to the top. Thus, the use of actionability is necessary in order to filter out these cases.

Proposed interventions generated by DrugAble need to be examined critically. Table 5 shows the top interventions for each of the top ten solutions. Solution 1, intervention 5 is

suggesting the use of an IL12 antibody to target IL10, even though it stands to reason that an IL12 antibody would act directly on IL12. Theoretically, this drug should have been filtered during search-space reduction because IL12 was not a prescribed target for this solution but is present in the model. However, upon examination of the drug-target database, "interleukin 12 antibody" was not annotated as modulating IL12. The first drug proposed by all interventions for Solution 5 is "lipid A derivatives." Lipid A in this case refers to the lipid component of lipopolysaccharide endotoxins common to gram negative bacteria[32]. These are being examined as potential adjuvants in vaccines because of their effect on immune modulators[33]. Rosmarinic acid, present in all of the top interventions for both Solutions 1 and 2, has gained recent attention for its potential ability to combat neurodegenerative diseases such as Alzheimers[34], but remains unverified as a useful pharmaceutical compound. Pulegone is a known carcinogen[35]. Oxprenolol and eplerenone are both anti-hypertensives. These drugs are commonly known as blood-thinners. Coagulation and related processes are known to be functions of the immune system, regulated by cytokines[36,37]. IC14 is a clinical trial-withdrawn CD14 antibody shown to inhibit TNF, IL6, and IL10, originally intended for the treatment of Amyotrophic Lateral Sclerosis (ALS)[38]. Seletalisib is an investigational P13K inhibitor being examined for the treatment of immune inflammatory diseases[39]. Taurolidine is an antimicrobial compound also being examined for the treatment of cancers[40], though the mechanisms for this are varied, not fully understood, and possibly hepatotoxic[40,41], despite the low number of annotated interactions. There are various other suggested drugs, many of which are monoclonal antibodies and other anti-cancer drugs. Most of these chemotherapeutics act through cytotoxic mechanisms, which are not ideal for treatment of non-cancers. These findings further reinforce the need for robust drug-target datasets free from contaminating interactions or drugs known to have undesired effects or mechanisms of action.

Using the number of off-target effects as a ranking for proposed interventions should be done with caution. These numbers will be as accurate as the database given to DrugAble. Any intervention with an exceptionally low number of off-target interactions, especially those with 0, warrants further analysis because it is rare that a drug is truly without off-target effects. The level of error in intervention generation is tolerable for two reasons. Firstly, it is often the case that multiple interventions are generated for a single solution, allowing experts to filter out incorrect combinations that may slip through. Furthermore, no proposed intervention would ever be implemented in a clinical trial without thorough examination by pharmacologists. DrugAble is not attempting to circumvent the need for expert examination, but rather to provide these same experts with tools to guide their work.

There are areas of improvement for DrugAble. Firstly, while drug-target interactions are accounted for by the actionability score, drug-drug interactions are not considered during the intervention generation phase. Future implementations of DrugAble may benefit from utilizing such knowledge when available, so as to avoid suggesting pharmaceutical interventions with known, adverse drug-drug interactions. Secondly, DrugAble may benefit from more straightforward tuning of actionability. The effects of τ and γ as tuning parameters are difficult to visualize, and there are currently no tuning parameters implemented for the repurposability or target scores. Thirdly, DrugAble is highly sensitive to poor data quality. Because it has no choice but to trust the drug-target information it is provided with, spurious results can occur when datasets have not been well-curated. Improvement to the specificity of queries made to DTQuery may improve mined dataset quality, but in any case, it is important to ensure that DrugAble is provided with robust datasets and that results are examined critically. Fifthly, while the filter-rank-filter intervention generation method is a clever approach to a complex problem, there is potential that

it will miss valid pharmaceutical interventions. It is also by far the slowest step in the DrugAble pipeline. Improving the speed and accuracy of intervention generation is therefore important. Finally, actionability scores are relative. It is not currently possible to compare actionability scores across models of disease with differing targets and drug-target datasets. If a method to generalize actionability scores across models and drug-target datasets were developed, it would become possible to assess the relative complexity of different models against one another. It would also be possible to assess the robustness of different drug-target datasets. Improving upon DrugAble in these ways would only enhance its usefulness and accuracy.

Despite the many areas for improvement, the current implementation of DrugAble is a marked improvement over current methods to design and assess multidrug interventions, which are few. Data and text mining methods for detecting off-target interactions have been previously applied by researchers such as Yera, but this data has not to our knowledge been applied toward the development of multidrug interventions. Computational prediction of drug-target effects is a major area of research in pharmacology[21–23,42,43], but predicted and putative effects are most-often used to guide traditional drug-development approaches. These new methods support the move away from traditional drug-development approaches and towards modern approaches that may prove to be more robust in treating complex illnesses.

An additional component score has been developed in order to further improve actionability, referred to as the *inter-model adjusted confidence score* (M_a). M_a is only implemented when multiple models, utilizing the same targets, are available to an analysis and have generated solutions. M_a accounts for the possibility that a given candidate solution may be generated by multiple models, and for the varying robustness and efficiency of those models. Robustness and efficiency are two measures generated simultaneously with each candidate

solution. Robustness *(r)* is the number of times that an MIS candidate succeeds in achieving a steady-state transition under a given number of simulations and efficiency (*m*) is the number of transitions required to achieve the goal steady-state. A detailed explanation of both can be found in Sedghamiz *et al.*, 2019. The formal definition of M_a for a given candidate solution *s* is provided in Equation (6):

Equation 6:

$$M_a(s) = |S| * M(r) * M(m)$$

where *S* is the set of all models containing *s*, *M*(*r*) is the median robustness of the models in *S*, and *M(m)* is the median efficiency of the models in *S*. When multiple models are present, maximizing M_a is ideal. However, when only one model is used in the analysis, or the analysis is not based on model-derived solutions, M_a is defaulted to 1 and thus has no effect on the final scoring. M_a will allow actionability to account for the variability in the tuning between models, an important step towards a more generalizable and robust version of actionability.

While the applicability of DrugAble is here demonstrated on chronic illness, specifically ME / CFS, there is no reason that it cannot be applied to a wide range of other illnesses, both chronic and acute. Furthermore, while used here on outputs from Bio-ModelChecker, this is not a strict requirement. So long as inputs are formatted in the way that DrugAble expects them, data from a variety of sources can be used. Rather than inputting an MIS candidate, proposed sets of target-action pairs can be given to DrugAble just as any MIS candidate from Bio-ModelChecker would be. Finally, while DTQuery currently only supports access to Elsevier's proprietary biological knowledgebase, any drug-target dataset can be used so long as it is formatted as a drug-target-action triplicate table. The flexibility in the implementation of DrugAble is important

because it allows researchers using any number of disease-model development methods and with virtually any drug-target dataset to also utilize DrugAble.

Conclusions

The prevalence, severity, and stubbornness of chronic diseases necessitates innovative methods for treatment discovery and design. The long-standing pharmacological paradigm of single-disease, single-drug interventions is both ill-founded and ineffectual for the treatment of chronic illnesses. In order to shift this paradigm, methods that support the development of both multi-drug interventions and the use of repurposable drugs are necessary. DrugAble is a promising new avenue toward the ranking of competing target-action sets and multidrug therapy design that attempts to address these issues. Its ability to effectively discriminate between theoretical solutions and those that are clinically actionable based solely on the network-structure of the target sets and drug-target interaction data can speed up treatment design significantly. By proposing computationally valid pharmaceutical combinations, researchers can be guided toward the use of repurposable drugs instead of developing new ones, a costly and time-consuming procedure. By reducing the barriers to multidrug treatment design and implementation, new paths towards total remission of ME / CFS and a variety of other chronic illnesses have become available.

References

1. Bernell, S. & Howard, S. W. Use Your Words Carefully: What Is a Chronic Disease? *Front. Public Health* **4**, (2016).
2. McGlynn, E. A. *et al.* The quality of health care delivered to adults in the United States. *N. Engl. J. Med.* **348**, 2635–2645 (2003).
3. Mannino, D. M. COPD: epidemiology, prevalence, morbidity and mortality, and disease heterogeneity. *Chest* **121**, 121S-126S (2002).
4. Wilson, H. D., Starz, T. W., Robinson, J. P. & Turk, D. C. Heterogeneity within the fibromyalgia population: theoretical implications of variable tender point severity ratings. *J. Rheumatol.* **36**, 2795–2801 (2009).
5. Abbi, B. & Natelson, B. H. Is chronic fatigue syndrome the same illness as fibromyalgia: evaluating the 'single syndrome' hypothesis. *QJM Int. J. Med.* **106**, 3–9 (2013).
6. Avellaneda Fernández, A. *et al.* Chronic fatigue syndrome: aetiology, diagnosis and treatment. *BMC Psychiatry* **9**, S1 (2009).
7. Somvanshi, P. R. & Venkatesh, K. V. A conceptual review on systems biology in health and diseases: from biological networks to modern therapeutics. *Syst. Synth. Biol.* **8**, 99–116 (2014).
8. Spivey, A. Systems Biology: The Big Picture. *Environ. Health Perspect.* **112**, A938–A943 (2004).
9. Deisboeck, T. S. Personalizing medicine: a systems biology perspective. *Mol. Syst. Biol.* **5**, 249 (2009).
10. Bielekova, B., Vodovotz, Y., An, G. & Hallenbeck, J. How implementation of systems biology into clinical trials accelerates understanding of diseases. *Front. Neurol.* **5**, 102

(2014).

11. Fritsch, P. *et al.* Succumbing to the laws of attraction. *Syst. Biomed.* **1**, 179–194 (2013).
12. Craddock, T. J. A. *et al.* Achieving Remission in Gulf War Illness: A Simulation-Based Approach to Treatment Design. *PloS One* **10**, e0132774 (2015).
13. Craddock, T. J. A. *et al.* A Logic Model of Neuronal-Glial Interaction Suggests Altered Homeostatic Regulation in the Perpetuation of Neuroinflammation. *Front. Cell. Neurosci.* **12**, (2018).
14. Tashkin, D. P. *et al.* A 4-Year Trial of Tiotropium in Chronic Obstructive Pulmonary Disease. *N. Engl. J. Med.* **359**, 1543–1554 (2008).
15. Baraniuk, J. N., El-Amin, S., Corey, R., Rayhan, R. U. & Timbol, C. R. Carnosine Treatment for Gulf War Illness: A Randomized Controlled Trial. *Glob. J. Health Sci.* **5**, 69–81 (2013).
16. Criner, G. J. *et al.* Benralizumab for the Prevention of COPD Exacerbations. *N. Engl. J. Med.* (2019) doi:10.1056/NEJMoa1905248.
17. Jayasundara, K. *et al.* Estimating the clinical cost of drug development for orphan versus non-orphan drugs. *Orphanet J. Rare Dis.* **14**, (2019).
18. Andronis, C., Sharma, A., Virvilis, V., Deftereos, S. & Persidis, A. Literature mining, ontologies and information visualization for drug repurposing. *Brief. Bioinform.* **12**, 357–368 (2011).
19. Corsello, S. M. *et al.* The Drug Repurposing Hub: a next-generation drug library and information resource. *Nat. Med.* **23**, 405–408 (2017).
20. Reddy, A. S. & Zhang, S. Polypharmacology: drug discovery for the future. *Expert Rev. Clin. Pharmacol.* **6**, (2013).
21. Yera, E. R., Cleves, A. E. & Jain, A. N. Prediction of off-target drug effects through data

fusion. in *Biocomputing 2014* 160–171 (WORLD SCIENTIFIC, 2013). doi:10.1142/9789814583220_0016.

22. Chartier, M., Morency, L.-P., Zylber, M. I. & Najmanovich, R. J. Large-scale detection of drug off-targets: hypotheses for drug repurposing and understanding side-effects. *BMC Pharmacol. Toxicol.* **18**, 18 (2017).
23. Torres, N. B. & Altafini, C. Drug combinatorics and side effect estimation on the signed human drug-target network. *BMC Syst. Biol.* **10**, 74 (2016).
24. Bulatov, A. A., Krokhin, A. A. & Jeavons, P. Constraint Satisfaction Problems and Finite Algebras. in *Automata, Languages and Programming* (eds. Montanari, U., Rolim, J. D. P. & Welzl, E.) 272–282 (Springer, 2000). doi:10.1007/3-540-45022-X_24.
25. Friedman, K. J., Bateman, L., Bested, A. & Nahle, Z. Editorial: Advances in ME/CFS Research and Clinical Care. *Front. Pediatr.* **7**, (2019).
26. Sepúlveda, N., Carneiro, J., Lacerda, E. & Nacul, L. Myalgic Encephalomyelitis/Chronic Fatigue Syndrome as a Hyper-Regulated Immune System Driven by an Interplay Between Regulatory T Cells and Chronic Human Herpesvirus Infections. *Front. Immunol.* **10**, (2019).
27. Frontiers | Cellular Immune Function in Myalgic Encephalomyelitis/Chronic Fatigue Syndrome (ME/CFS) | Immunology. https://www.frontiersin.org/articles/10.3389/fimmu.2019.00796/full.
28. Morris, M. C. *et al.* Leveraging Prior Knowledge of Endocrine Immune Regulation in the Therapeutically Relevant Phenotyping of Women With Chronic Fatigue Syndrome. *Clin. Ther.* **41**, 656-674.e4 (2019).
29. Taneja, V. Sex Hormones Determine Immune Response. *Front. Immunol.* **9**, (2018).
30. Kerr, J. R. Epstein-Barr Virus Induced Gene-2 Upregulation Identifies a Particular Subtype

of Chronic Fatigue Syndrome/Myalgic Encephalomyelitis. *Front. Pediatr.* **7**, (2019).

31. Richman, S. *et al.* Pharmaceutical Interventions in Chronic Fatigue Syndrome: A Literature-based Commentary. *Clin. Ther.* **41**, 798–805 (2019).
32. Lodowska, J., Wolny, D., Weglarz, L. & Dzierzewicz, Z. [The structural diversity of lipid A from gram-negative bacteria]. *Postepy Hig. Med. Doswiadczalnej Online* **61**, 106–121 (2007).
33. Martin, M., Michalek, S. M. & Katz, J. Role of Innate Immune Factors in the Adjuvant Activity of Monophosphoryl Lipid A. *Infect. Immun.* **71**, 2498–2507 (2003).
34. Hase, T. *et al.* Rosmarinic acid suppresses Alzheimer's disease development by reducing amyloid β aggregation by increasing monoamine secretion. *Sci. Rep.* **9**, 1–13 (2019).
35. Jabba, S. V. & Jordt, S.-E. Risk Analysis for the Carcinogen Pulegone in Mint- and Menthol-Flavored e-Cigarettes and Smokeless Tobacco Products. *JAMA Intern. Med.* **179**, 1721–1723 (2019).
36. Antoniak, S. The coagulation system in host defense. *Res. Pract. Thromb. Haemost.* **2**, 549–557 (2018).
37. Burzynski, L. C. *et al.* The Coagulation and Immune Systems Are Directly Linked through the Activation of Interleukin-1α by Thrombin. *Immunity* **50**, 1033-1042.e6 (2019).
38. Axtelle, T. & Pribble, J. IC14, a CD14 specific monoclonal antibody, is a potential treatment for patients with severe sepsis. *J. Endotoxin Res.* **7**, 310–314 (2001).
39. Helmer, E. *et al.* First-in-human studies of seletalisib, an orally bioavailable small-molecule PI3Kδ inhibitor for the treatment of immune and inflammatory diseases. *Eur. J. Clin. Pharmacol.* **73**, 581–591 (2017).
40. Neary, P. M. *et al.* The Evolving Role of Taurolidine in Cancer Therapy. *Ann. Surg. Oncol.*

17, 1135–1143 (2010).

41. Fahrner, R. *et al.* Short-term treatment with taurolidine is associated with liver injury. *BMC Pharmacol. Toxicol.* **18**, (2017).
42. Takarabe, M., Kotera, M., Nishimura, Y., Goto, S. & Yamanishi, Y. Drug target prediction using adverse event report systems: a pharmacogenomic approach. *Bioinformatics* **28**, i611–i618 (2012).
43. Pliakos, K. & Vens, C. Drug-target interaction prediction with tree-ensemble learning and output space reconstruction. *BMC Bioinformatics* **21**, 49 (2020).