Predicting the Immune Response to Repurposed Drugs in Coronavirus-induced Cytokine Storm

Matthew C Morris
Center for Clinical Systems Biology
Rochester General Hospital
Rochester, NY, USA
matthew.morris@rochesterregional.org

Cole A Lyman
Center for Clinical Systems Biology
Rochester General Hospital
Rochester, NY, USA
cole.lyman@rochesterregional.org

Spencer Richman
Center for Clinical Systems Biology
Rochester General Hospital
Rochester, NY, USA
richmanspencer@gmail.com

Hong Bao Cao
Biology Solutions
Elsevier BV
Amsterdam, Netherlands
h.cao@elsevier.com

Chris Cheadle
Biology Solutions
Elsevier BV
Amsterdam, Netherlands
c.cheadle@elsevier.com

Gordon Broderick
Center for Clinical Systems Biology
Rochester General Hospital
Rochester, NY, USA
gordon.broderick@rochesterregional.org

Abstract— **A significant cause of morbidity in COVID-19 infected patients admitted to the hospital is a severe dysregulated inflammatory response characterized as a cytokine storm, a key component of acute respiratory distress syndrome (ARDS). Here we have assembled a basic immune regulatory model from a list of 19 immune mediators with reported involvement in cytokine storm. Automated text-mining of over 2,500 full text journal publications using the MedScan natural language processing (NLP) engine identified 112 documented regulatory interactions coordinating the dynamic response of this network. This same text mining highlighted reported bi-directional associations between Coronavirus infection and a broad set of immune mediators producing a complex feedback pattern of host-pathogen interaction. Decisional logic parameters supporting the network's dynamic response were identified such that observed responses to SARS-CoV infection in an *in vitro* system of Calu3 human lung adenocarcinoma cells could be accurately predicted. Of the 19 competing models, 2 supported a dominant inactive immune resting state, with a predicted onset of cytokine storm in 63% and 26% of simulated infections respectively. Discrete event simulation based on the latter suggest that some repurposing strategies might outperform popular use of hydroxychloroquine as a companion to anti-viral therapy.**

Keywords—Coronavirus, rapid prototyping, discrete logic, simulation, natural language processing

I. Introduction

Coronaviruses infecting humans (hCoV) are an emerging family of viruses increasingly responsible for serious disease. Prior to the COVID-19 pandemic, major coronavirus outbreaks occurred in 2002 (SARS) and 2012 (MERS) [1], with case fatality rates of approximately 10% for SARS [2] and 35% for MERS [3]. As the cause of the more recent outbreak, MERS-CoV has been the focus of more active vaccine development efforts, though as of 2019 none of these had advanced beyond early clinical trials in humans [4]. There is currently no vaccine available for COVID-19, though an experimental recombinant vaccine administered via microneedles has been shown to elicit neutralizing antibodies in mice [5] and an adenovirus-vectored vaccine has completed phase I trials in human subjects in China [6]. Both SARS and MERS infect airway epithelial cells of the lower respiratory tract, potentially causing acute lung injury (ALI) and progressing to acute respiratory distress syndrome (ARDS) in the most serious cases [7],[8]. Immunopathology is a major contributor to the high morbidity and mortality rates from coronavirus infection [1], with the most severe disease generally occurring in immunocompromised or elderly patients, or those with comorbidities [2]. Indeed, immunopathology appears to be the determining factor of the course of disease, as fatal hCoV infection is usually marked by progression to ALI and ARDS after peak viral replication [2].

These viruses are known to possess sophisticated mechanisms for immune evasion, especially by suppressing IFN response [2]. Infection is marked by delayed but elevated expression of pro-inflammatory cytokines, especially Type I IFN and interferon-stimulated genes. Levels of these cytokines have been found to be correlated with disease severity [2]. Early in infection, coronaviruses delay IFN response, which enables rapid replication in airway epithelial cells [2]. Dysregulation of Type I IFN production was found to be a major determinant of the course of SARS-CoV infection in mice. Early IFN production results in rapid viral clearance, preventing later immunopathology, while delayed IFN production promoted viral replication and led to ARDS. Strikingly, IFN-KO mice experienced heightened viral replication but did not progress to ARDS [9], demonstrating that immunopathology and not viral replication is necessary for fatal disease. While efforts continue towards developing hCoV-specific antiviral drugs by blocking viral proteins or suppressing host elements essential for

DOI 10.1109/BIBE50027.2020.00080

replication, but none of these have yet proven unequivocally successful [10]. Of equal importance and of specific relevance to this work, the deadly acute immune response to hCoV infection has been treated with limited success by corticosteroids or interferon, but precise timing is critical [2], likely making these options infeasible for general use. Indeed, the central role of immunopathology in severe hCoV disease suggests that attacking the virus directly may not be the optimal method of treatment at all. Modulation of host immunity, especially the IFN response, has been proposed as a more feasible avenue for reducing the severity and morbidity of hCoV disease [2],[10].

The COVID-19 outbreak has caused a surge of interest in possible use of the antimalarial drug quinine and its derivatives for hCoV treatment. Chloroquine is known to suppress viral replication in vitro [11], [12], but failed to protect human patients from infection by influenza or dengue viruses in clinical trials [11]. Again, the role of immunopathology must not be overlooked. Quinine derivatives have shown immunomodulatory effects, especially with respect to IFN, particularly in the context of response to intracellular RNA via TLR7 [13]. This last point is of special interest for possible applications in hCoV treatment. If quinine derivatives are effective in hCoV treatment, their immunomodulatory effects are likely to be at least as important as their suppression of viral replication. Immunomodulation may constitute the best means of averting or alleviating immunopathology in hCoV infection. Predicting appropriate targets for this purpose is made challenging not only by the sheer complexity of immune regulatory programming but also by virtue of the relative scarcity of molecular and cellular data describing onset and progression of COVID-induced ARDS and related cytokine storm. Moreover, available evidence suggests that COVID-19 related ARDS may represent a specific variant thereof [14].

Emerging work on early predictive markers of ARDS onset, specifically in the context of COVID-19 infection, have so far pointed to changes in relative abundance of specific lymphocyte subsets as a possible indicator of severity and prognosis of pneumonia in these patients [15]. However, under the current front-line circumstances, studies such as these are still relatively few in number, typically survey only a small subset of select markers and consist of subjects recruited according to a fairly broad and varying range of entry criteria. As such the data collected is often poorly suited for conventional statistical modeling and these analyses can be especially prone to producing spurious results [16]. Moreover, although many studies are longitudinal in design, markers continue to be examined independently and without formal consideration of the underlying regulatory dynamics. Most importantly, these studies remain purely empirical and hence do not implicitly consider known and validated immune response mechanisms leaving the COVID-specific etiology of ARDS and the design interventional approaches to disrupt the related cytokine storm difficult to infer with great confidence.

In this work we apply a literature-informed approach that is inherently robust to sparse and incomplete data to investigate the onset and progression of cytokine storm response to hCoV infection and to simulate the anticipated effects of applying known drugs to disrupting and remediating this often-fatal complication. We assemble a mechanistic regulatory network of immune signaling using the broad-scale automated text-mining over 2,500 journal publications and require the latter recover the inflammatory response kinetics of Calu3 human lung adenocarcinoma cells challenged transiently *in vitro* with SARS-CoV [17]. Recovery of this *in vitro* time course was supported by 19 competing models however only 2 also supported a dominant stable attractor corresponding to an inactive immune resting state. Using these models, we found that the predicted regulatory effects of hydroxychloroquine did in fact partially destabilize the basin of attraction corresponding to persistent cytokine storm. However, simulations also found that Ruxolitinib, another drug under clinical investigation, delivered an even more substantial down-regulation of cytokine storm. Importantly, when used in conjunction with an idealized anti-viral, the latter delivered lasting stable resolution and recovery of a relatively inactive immune resting state. In comparison, quinine derivatives though helpful in provided transient benefits were not predicted to deliver a lasting resolution even when combined with an anti-viral.

II. Methods

A. Assembling a prototype literature-informed network

A basic immune regulatory network was assembled from molecular markers with documented involvement in infectious pneumonia as extracted from the Elsevier Knowledge Graph (Elsevier, Amsterdam) using the Pathway Studio* suite of software tools [18]. The network consists of 19 immune mediators linked by 112 regulatory actions (edges), extracted from a total of 2,653 published references (with a median of 7 references per edge). In this network, any given immune mediator is itself regulated by up to 10 upstream mediators (maximum in-degree). This regulatory input space corresponds to an exhaustive truth table (K matrix) with a total of 3,390 transition state entries. The network contains only 1 sink or output node (CD80) and no purely exogenous source node. Coronavirus is represented in Elsevier Knowledge Graph database as a "disease", with regulatory actions on immune mediators CCL5, CD200R1, CD40, CD80, CD86, CSF3, CTSB, CTSL, CXCL10, CXCL2, FGL2, IFNL1, STAT2, TNF, STAT1, NFKB2, NFKB1, and IFNG. Similarly, Coronavirus is acted upon directly or indirectly by the host immune network through feedback from CD200R1, STAT1, NFKB2, NFKB1, IFNG. In addition, the self-perpetuating nature of an active infection is represented here as a positive feedback of coronavirus onto itself. Though other immune modulators and cell populations are crucial for the resolution of infection, this minimal set was judged to be representative of the core host response signaling mechanisms and a sufficient initial foundation for validation against observed experimental data and generally expected stable homeostatic behaviors (Fig. 1).

B. Inferring novel regulatory actions

This first immune signaling circuit was constructed using known immune regulatory interactions documented in the literature and extracted from the Elsevier Knowledge Graph. However, there are still many gaps in our understanding of immune signaling and this is especially true of host-pathogen interactions involving novel pathogens such as members of the Coronavirus family. Many approaches for network discovery

US Department of Defense CDMRP, Elsevier BV

have been proposed many of which rely on reverse engineering from relatively large amounts of data, in particular informative data such as structured perturbation experiments [16]. However, in situations such as this where experimental data is sparse and collected under a narrow range of conditions, these data-driven methods typically perform poorly, leading to excessively high false positive rates. Under such conditions, methods rooted in a graph theoretical analysis of network structure can serve to identify connectivity patterns which deviate from those known to be characteristic of biological networks, for example highly connected modules and hub-like architecture. Here we apply one such structure-based approach proposed by Guimerà and Sales-Pardo (2009) [19] based on stochastic block models whereby connections between nodes are typically more abundant within modules than between modules. This property is used here to estimate the probability that a novel and currently absent interaction might exist given the structure of the newly augmented network.

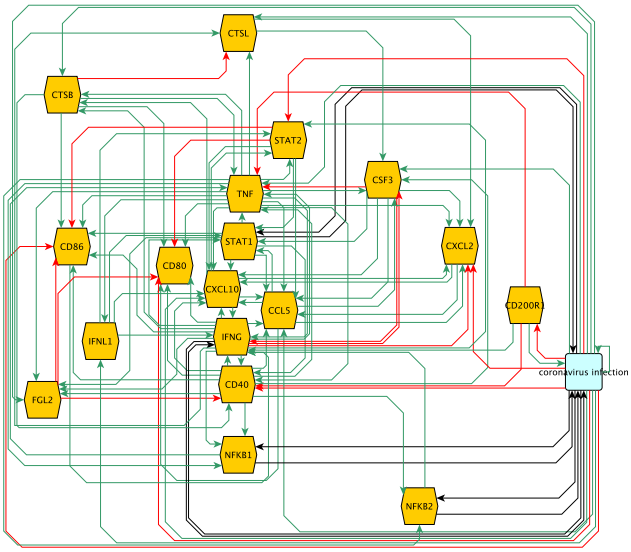

Fig. 1. *A model immune regulatory circuit*. In causal immune regulatory circuit assembled from the automated text mining of 2,653 journal publications, 18 immune mediators and the coronavirus pathogen were linked by 112 documented activating (green edges) and inactivating (red edges) regulatory actions. A secondary analysis of network structure suggested the addition of hypothetical regulatory connections between coronavirus and 4 other network entities (black edges), for which both direction and effect were unknown.

While signaling among elements of the host immune system have been well studied, this is much less true of pathogen-host interactions involving COVID-19, making these of special interest. In the case of this initial network consisting of 112 documented regulatory interactions, we found we found possible novel interactions linking 4 host immune mediators with coronavirus with a probability of occurrence, exceeding 0.50 (Table I). As the algorithm predicting interaction reliability will infer the presence or absence of undirected edges, these candidate regulatory actions were incorporated into the original network in a bi-directional fashion to account for both possible cases. Hypothetical edges were retained in the final network only if they improved adherence to the data to a greater degree than they increased the complexity of the model as detailed below. This augmented regulatory circuit is shown in Fig. 1.

TABLE I. HYPOTHETICAL MISSING INTERACTIONS INFERRED FROM THE ANALYSIS OF NETWORK STRUCTURE.

Edge	Probability
coronavirus infection – STAT1	0.72
coronavirus infection – NFKB1	0.65
coronavirus infection – NFKB2	0.65
coronavirus infection - IFNG	0.61

C. *Identifying plausible sets of decisional logic parameters*

The expression of networked host-pathogen immune elements is described qualitatively (e.g. low, nominal and high) and evolves in time as dictated by the control actions of neighboring entities (e.g. decrease, increase, or remain unchanged). These control actions are activated in a threshold- and context-specific manner. For example, in Fig. 2 entity A will inhibit entity B only when it is expressed at greater than nominal levels (activation threshold) and will result in a decrease of entity B expression only when it is acting in the absence of the competing strong promoter C (context-specific transition logic). Values are selected for these logic parameters at every node such that the predicted dynamic behaviors supported by the network must allow for the accurate recovery of all available experimentally observed responses and expected resting states. The sheer number of permutations in these parameter values gives rise to a combinatorial explosion of potential competing models. To address this, our group has formulated this as an efficient multi-objective Constraint Satisfaction Problem (CSP) [20] which is resolved here using the Chuffed solver [21]. Parameters describing each regulatory interaction are constrained to comply with the direction and mode of action (activation or inactivation) reported in the literature. Postulated stationary states must also be reproduced exactly (hard constraints) as the solver minimizes error in the prediction of experimentally observed transient states, while also minimizing the model's structural (presence or absence of a regulatory interaction) and regulatory (required transition logic resolution) complexity. Prediction error is calculated here as the absolute difference between the simulated output and the experimental time course data. The structural efficiency is defined as the number of regulatory interactions and regulatory efficiency as the minimal number of distinct activation threshold values and instances of distinct regulatory outcomes for a node among all possible combinations of active input. Since the available data are never sufficient to fully constrain logical parameters in all contexts, multiple parameterizations can support equally low values for the objective function. These parameterizations support the available data equally well, but predict regulatory dynamics which may differ quite considerably from one another in contexts unconstrained by input data.

US Department of Defense CDMRP, Elsevier BV

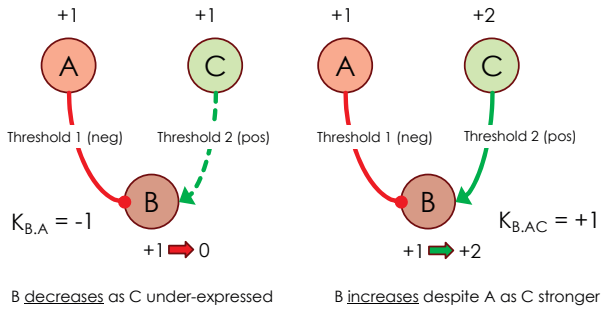

Fig. 2. *Discrete logic decisional parameters.* In a basic circuit a node B can be downregulated by node A and upregulated by node C. In the left panel, when A is expressed at a state +1 in excess of its activation threshold whereas C is expressed at state +1 below its activation threshold, then node B is downregulated by node A acting alone ($K_{B,A}$ = -1). In the right panel, when A is expressed at a state +1 in excess of its activation threshold but C is further expressed at state +2 above its activation threshold, then node B is regulated by both nodes. The net decisional weight of nodes A and C acting on B ($K_{B,AC}$ =+1) is such that the net regulatory action is to upregulate B.

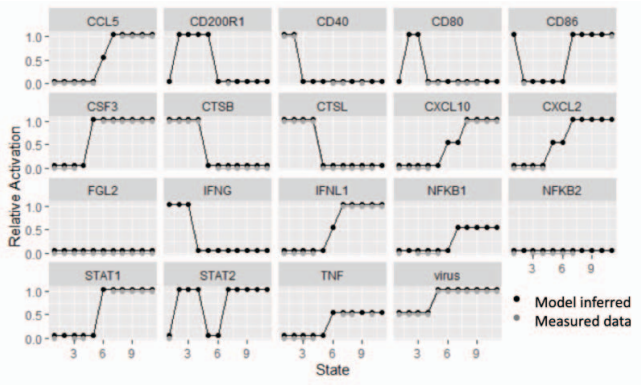

Fig. 3. *Discrete state in vitro response to infection.* Experimental data (grey dots) describing in vitro response over 72 hours of human Calu-3 cells to SARS-CoV-infection discretized by variational Bayes Gaussian (VBG) clustering. Points show the cluster medians for each marker at the indicated timepoint where measurements could be assigned to a discrete value with at least 90% confidence (57.4% of the data). Model-predicted levels are superimposed (black dots).

D. Steady States and Basins of Attraction

Given an initial state of the network, the next state towards which the network should evolve (or image) can be predicted by applying one of possibly several competing sets of decisional logic parameters. This next target state or network image can be applied to a single randomly selected node (asynchronous updating) or uniformly and simultaneously across all nodes in the network (synchronous updating). The latter case of synchronous updating produces a much more compact state transition graph (STG) and delivers identical stationary points or stable resting states, albeit at the exclusion of more complex stable oscillatory behaviors. Here we focus specifically on stationary non-oscillatory resting states and apply a synchronous update scheme to monotonically update (increase or decrease by 1) all network nodes where the next state predicted by the decisional logic or their image differs from their current state. Let a steady state or equilibrium be defined as an instance where the network image is equal to the current state of the network indicating that no nodes are eligible to be updated [22]. Furthermore, let a basin of attraction be a set of states associated with a steady state, such that when a simulation is started from an element of the basin of attraction it leads to the steady state within a given number of transitions.

In this work, all of the steady states for each model are identified using a constraint-based optimization formulated in MiniZinc [23]. The basins of attraction were identified using a stochastic sampling method that simulated from 100,000 random start states and identified how frequently a steady state was reached within 50 transitions. The relative frequency that a steady state is reached from a random state is an approximation for the size of the basin of attraction, where the larger the basin of attraction the more frequently the associated steady state will be reached.

E. Discrete state transformation of experimental data

Existing a previous *in vitro* study of human Calu-3 lung adenocarcinoma cells infected by SARS-CoV over 72 hours was retrieved from the Gene Expression Omnibus (GEO accession number GSE33267). Gene expression of network entities in SARS-CoV-infected samples was expressed as a fold change relative to mock-infected cells. These continuous measurements were then converted to discrete values by unsupervised clustering using variational Bayesian Gaussian methods [24],[25]. Each entity was discretized independently of the others using the full range of values from both mock-infected and SARS-CoV-infected samples over the whole time course (grey dots, Fig. 3). The number of distinct activation levels available was determined by the algorithm, with an upper limit of 4. Where measurements could not be assigned to a discrete value with at least 90% confidence, they were labelled as missing data. This step resulted in the retention of 57.4% of available data points. The median discretized values for each infection condition at each timepoint were then used as a summary, yielding the final discretized infection trajectory for parameterization (Figure 3). Based on measurements of viral titer over the course of the infection performed by Sims et al. (2013)[17], coronavirus was stipulated to begin at an active level before increasing to its maximum at approximately 24 HPI, where it persisted for the remainder of the time course, but inflammatory responses continued to intensify throughout the time course even after viral titers had peaked. The model was additionally constrained to capture the instance of complete viral clearance by prohibiting predictions of coronavirus spontaneously re-emerging after reaching zero titer.

III. RESULTS

A. Recovering in vitro experimental observations

Decisional logic parameters dictating the dynamic behavior of the model were discovered using a constraint satisfaction framework developed by our group [20]. The Chuffed solver [21]was applied to the regulatory circuit in Fig. 1 and directed to minimize both departure from the experimentally observed *in vitro* response trajectories and network complexity. With these settings, 19 unique models were identified with <5% departure from the experimental data. Of these, 11 models matched the available data exactly. The median predicted response trajectories from all 19 models are depicted in Fig. 3, showing

adherence to the available experimental data as well as predicted activation values for entities where measurements were unavailable. Of the hypothetical interactions included in the model, a downregulation of coronavirus by IFNG was retained in 17 of the 19 models, indicating with a high degree of consensus that this host-pathogen interaction was necessary to accurately recover experimentally observed behavior. Of the 3390 K parameters describing the incremental state transition rules, 540 were invariant across all 19 models, but there was substantial variation among remainder (median Manhattan distance 1274 out of a maximum 5402, or 23.6%). The median predicted activation states for each immune mediator across all 19 of the top-performing models nonetheless shows close adherence between models to the available measured response trajectories. A detailed examination of the sources of the remaining error for each model can be found in Figure 4.

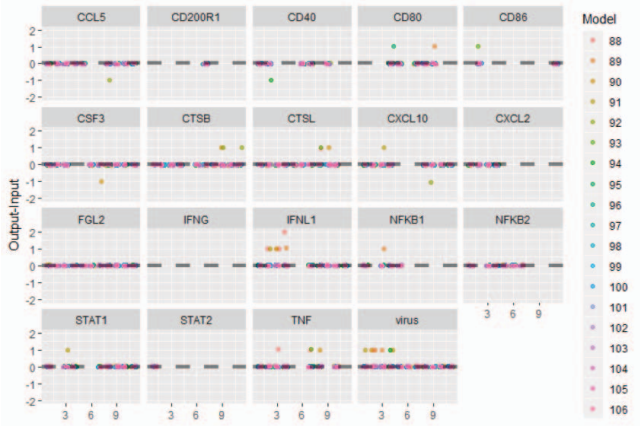

Figure 4. *Recovery of measured immune marker expression*. Departure of predicted discrete activation state for each immune marker from measured values for each of the top 19 models. Dashed lines represent exact match of predicted to observed values.

B. Comparing to expected stability and incidence

These 19 models constitute a family of competing hypotheses for the progression of coronavirus infection in this context as informed by the available input data. To assess their ability to support expected regulatory regimes, each model was simulated 100,000 times from random starting positions to map the available basins of attraction (Figure 5). These random simulations were necessary as the total state transition graph comprised over 10^7 states, which was too large to examine exhaustively. Simulations were conducted from randomly chosen start states, then again with coronavirus titer set to its maximum (2) or not detectable (0). Each simulation was performed over a horizon of 50 transitions or until reaching an attractor, defined as any state for s(t) which its successor s(t+1) was equal to s(t). Multidimensional scaling was used to compare identified attractors 1) to each other, 2) to a hypothetical inactive immune resting state, and 3) to the 72 HPI timepoint predicted by each model and stipulated here to represent cytokine storm.

In simulations conducted from initial states where coronavirus infection was absent (s_{cv}(all t)=0), the putative cytokine storm states were never reached by any model. Most models supported multiple basins of attraction not far removed from the hypothetical inactive immune resting state however these basins were typically quite small, capturing only a small minority of total simulations (mean of 12% <<50%) (Fig. 5). Only 2 of the 19 models (models 13 and 18) supported stable states in close proximity (Manhattan distance of 6) to immune inactivation and towards which a majority of simulations came to rest in the absence of infection (73.0 and 77.7% of occurrences respectively). This is consistent with the intuitive expectation that a state of relative immune inactivation should predominate in the absence of active infection. Model 18 additionally supported a steady state with a Manhattan distance from the inactive immune state of only 1, which was reached in 17.6% of simulations conducted in the absence of coronavirus infection.

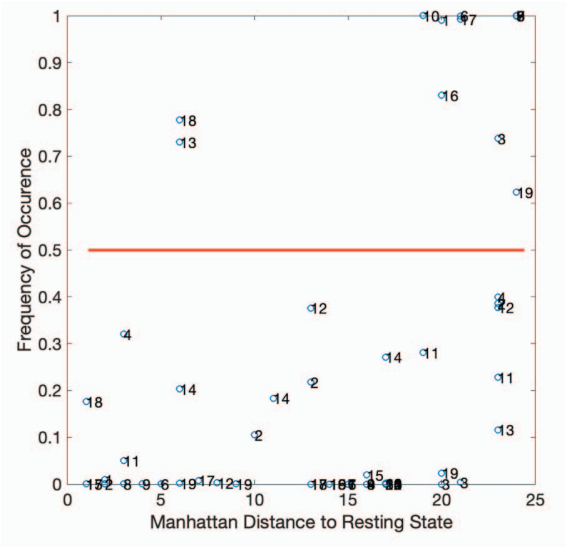

Figure 5. *Recovery of an inactive immune resting state*. Departure from an idealized inactive immune resting state of stable attractors identified for each of the 19 candidate models in the absence of infection. Simulations conducted with models 13 and 18 come to rest at states closest (Manhattan distance of 6) to broad immune activation in a majority of simulations (72 and 78% respectively).

When initiated from states with an active coronavirus infection (s_{cv}(all t)=2), virtually all models most frequently migrated to their predicted cytokine storm terminal states. One model predicted spontaneous resolution of infection, but at a cost of substantial persistent immune activation. Both models 13 and 18 escaped to a persistent immune activation proximal to cytokine storm in the presence of active coronavirus infection in 63.2% and 26.4% of occurrences respectively.

TABLE II. DOCUMENTED TARGETS AND MODE OF ACTION FOR BROADLY STUDIED HYDORXYCHLOROQUINE AND THE MUCH LESS STUDIED RUXOLITINIB.

Drug	Documented Actions
Hydroxychloroquine	CD80 positive, CD86 positive, CTSB negative, CTSL negative, CXCL10 positive, IFNG negative, TNF negative
Ruxolitinib	CXCL10 negative, IFNG negative, STAT1 negative, TNF negative

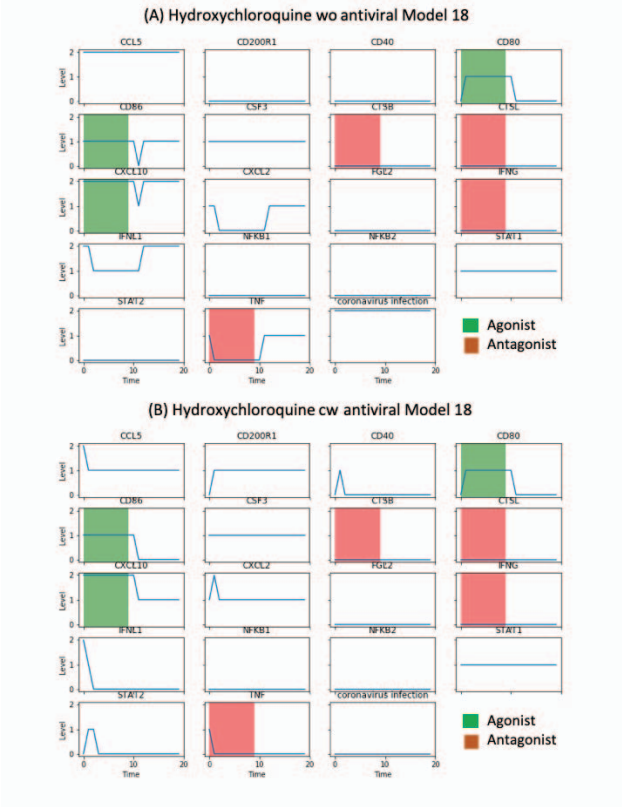

Fig. 6. *Simulated use of hydroxychloroquine*. The application of hydroxychloroquine agonist (green bars) and antagonist (red bars) effects to an established cytokine storm is simulated with a maximal viral load (A) and with the complete clearance of virus through concurrent use of an idealized antiviral (B).

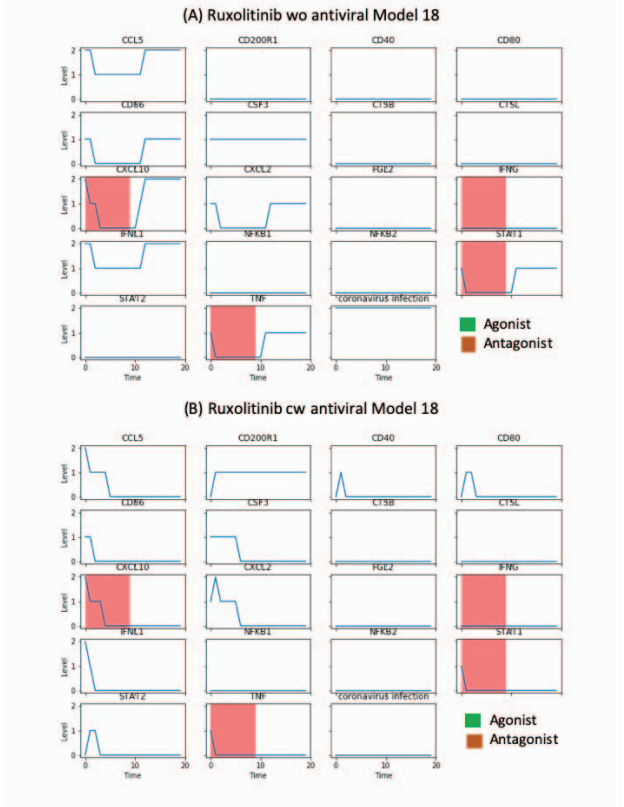

Fig. 7. *Simulated use of ruxolitinib*. The application of ruxolitinib's antagonist (red bars) effects to an established cytokine storm is simulated with a maximal viral load (A) and with the complete clearance of virus through concurrent use of an idealized anti-viral (B).

C. Simulating strategies for disrupting cytokine storm

At the time of this writing, a total of 706 interventional clinical trials were registered, active and/ or recruiting globally to study COVID-19 induced severe acute respiratory syndrome (SARS-CoV). Of these 149 involved potential therapeutic applications of quinine and its derivatives, with 110 of these assessing the use of hydroxychloroquine and to a lesser extent chloroquine (39 studies). In contrast drugs such as dexamethasone (8 studies) and ruxolitinib (6 studies) are attracting more recent interest and have not yet been extensively studied for use in treating COVID-19 pneumonia. The documented actions of hydroxychloroquine, chloroquine and quinine on immune mediators in our model regulatory circuit were extracted from the Elsevier Knowledge Graph (Elsevier, Amsterdam) and are listed in Table II. Note that while quinine derivatives have been reported to suppress viral replication [11],[12], they have not been demonstrated to specifically inhibit coronavirus, and they are therefore not represented as directly impacting coronavirus levels. For the purposes of this analysis we also chose to simulate the effects of ruxolitinib, a Janus-associated kinase (JAK1/2) inhibitor currently used to treat lupus, as a more novel and less studied comparator intervention [26]. Once again, the molecular targets and mode of action for ruxolitinib were extracted from the Elsevier Knowledge Graph database and are shown in Table II as downregulating the expression of CXCL10, IFNG, STAT1 and TNF.

Recall that model 18 supported an onset of stable immune hyperactivation in 26% of the simulated infections, a number which closely resembles the complication rate for COVID-19 infections in the general population [27]. As such, we used the latter to predict the potential efficacy of both hydroxychloroquine and ruxolitinib. Initializing the simulation from an established immune hyperactivation resembling cytokine storm, we applied each drug with and without an idealized anti-viral agent, mimicking their frequent use as companion drugs in current clinical protocols. Synchronous update was used to simulate the evolution of the immune response over a horizon of 20 time steps with these drugs applied at the outset and maintained for the first 10 time steps.

Simulations involving hydroxychloroquine in the context of a full viral load (Fig. 6A) suggested only a limited ability to modulate inflammation. Initially, a Manhattan distance of 11 separates the immune expression profile during a putative cytokine storm from that of an idealized inactive immune resting state (all zero). The simulated use of hydroxychloroquine transiently reduced this at best to a Manhattan distance of 7 from quiescent immune state. However, if the viral load was reduced by concurrent use of an idealized antiviral, immune activation

was reduced to a Manhattan distance of 6 to the target quiescent state and persisted after cessation of treatment.

In the case of ruxolitinib, applied in the absence of an antiviral at full viral load, model 18 predicted that the immune activation profile would be brought to within a Manhattan distance of 3 from the target idealized immune inactivation state. If the virus was simultaneously cleared through concurrent use of an antiviral then ruxolitinib was predicted to resolve immune hyperactivation to within a Manhattan distance of 1 from broad immune quiescence, or the same resting state attractor identified using stochastic searches. This state involved persistent intermediate activation of CD200R1 with all other markers at their minimum levels, and remained stable even after cessation of treatment.

IV. Discussion

The high morbidity and mortality associated with human coronavirus infections is due largely to the detrimental patterns of dysregulated immune activation triggered by these viruses. By constraining a family of discrete logical models to support an existing dynamic trajectory of immune response to SARS-CoV *in vitro* over time, we identified regulatory regimes which could be easily driven into a cytokine storm by coronavirus infection. The system could be dislodged from these states by simulated hydroxychloroquine treatment, though a return to stable rest remained unlikely. Simulations of a more novel and much less studied intervention involving the drug ruxolitinib compared very favorably, reducing the activation of immune markers to a greater degree. With successful antiviral treatment concurrently reducing viral load to non-detectable levels, the latter was predicted to support virtually full and lasting resolution of cytokine storm.

The central caveat to this work is that these regulatory logic models were parameterized using data from an in vitro study of immortalized lung epithelial cells. While SARS-CoV is known to infect such cells in human patients [28] and the Calu3 model studied here has been frequently employed [29], such systems necessarily contain only one cell type. The many specialized immune cells mobilized against an infection are therefore not present, and cannot be modeled with these data. This is a major reason for the inability of the models to resolve coronavirus infection, and for the emergence of cytokine storm as a stable steady state. The immune response studied here is governed entirely by autocrine and paracrine signaling among Calu3 cells.

With this caveat in mind, our simulations were successful in recapitulating the available data, suggesting a reasonably accurate capture of the regulatory dynamics at work. Since the reference data were not sufficient to exhaustively constrain the entirety of the dynamic parameters under all possible contexts, we were left with a family of competing models, all equally consistent with the available data. By simulating the predicted behavior of these models under different conditions, we found that they shared a tendency to be driven into a state of stable immune hyperactivation by coronavirus infection. Since only two of the competing models supported a state of relative immune quiescence as a stable attractor, and one of these aligned particularly well with observed complication rates, we focused further simulations on this model, finding that the predicted effects of hydroxychloroquine and ruxolitinib were both sufficient to destabilize cytokine storm, preventing the system from arriving there during coronavirus infection.

Together, our results suggest a potential explanation for the thus-far anecdotal observations of hydroxychloroquine benefits in coronavirus treatment and encourage the continued study of ruxolitinib. Once dislodged from cytokine storm, the complete *in vivo* immune system may be better enabled to return to its healthy rest state—in effect "buying time" for the viral load to peak and subside without dangerous overactivation and progression to ARDS. To reliably avert the onset of cytokine storm and permit clearance of the precipitating infection, the targets and timing of pharmaceutical interventions must be precisely chosen based on the regulatory logic of the system as a whole.

Acknowledgment

This work was supported by Rochester Regional Health in conjunction with the US Department of Defense Congressionally Directed Medical Research Programs (CDMRP) http://cdmrp.army.mil/ under Peer Reviewed Medical Research Program (PRMRP) award W81XWH1910804 (Broderick - PI; Sethi - Partnering PI), as well as by Elsevier AB (Amsterdam). The authors would also like to thank Drs. Chris Cheadle, Hongbao Cao, Dicle Hasdemir, Anton Yuriev, Philipp Anokhin, Umesh Nandal and Thibault Geoui of Elsevier BV for their technical assistance and many helpful discussions.

MANDATORY DISCLAIMER

The opinions and assertions contained herein are the private views of the authors and are not to be construed as official or as reflecting the views of the Department of Defense.

References

[1] Gralinski LE, Baric RS. Molecular pathology of emerging coronavirus infections. J Pathol. 2015;235(2):185-195. doi:10.1002/path.4454

[2] Channappanavar R, Perlman S. Pathogenic human coronavirus infections: causes and consequences of cytokine storm and immunopathology. Semin Immunopathol. 2017;39(5):529-539. doi:10.1007/s00281-017-0629-x.

[3] Chafekar A, Fielding BC. MERS-CoV: Understanding the Latest Human Coronavirus Threat. Viruses. 2018;10(2):93. Published 2018 Feb 24. doi:10.3390/v10020093.

[4] Schindewolf, C.; Menachery, V.D. Middle East Respiratory Syndrome Vaccine Candidates: Cautious Optimism. Viruses 2019, 11, 74. doi: 10.3390/v11010074

[5] Kim E, Erdos G, Huang S, et al. Microneedle array delivered recombinant coronavirus vaccines: Immunogenicity and rapid translational development. EBioMedicine. 2020;55: 102743. doi:10.1016/j.ebiom.2020.102743.

[6] Zhu FC, Li YH, Guan XH, et al. Safety, tolerability, and immunogenicity of a recombinant adenovirus type-5 vectored COVID-19 vaccine: a dose-escalation, open-label, non-randomised, first-in-human trial. Lancet. 2020;395(10240):1845-1854. doi:10.1016/S0140-6736(20)31208-3

[7] Chafekar, A.; Fielding, B.C. MERS-CoV: Understanding the Latest Human Coronavirus Threat. Viruses 2018, 10, 93. doi : 10.3390/v10020093

[8] Weiss SR, Leibowitz JL. Coronavirus pathogenesis. Adv Virus Res. 2011;81:85-164. doi:10.1016/B978-0-12-385885-6.00009-2.

[9] Channappanavar R, Fehr AR, Vijay R, et al. Dysregulated Type I Interferon and Inflammatory Monocyte-Macrophage Responses Cause

Lethal Pneumonia in SARS-CoV-Infected Mice. Cell Host Microbe. 2016;19(2):181-193. doi:10.1016/j.chom.2016.01.007

[10] Totura AL, Bavari S. Broad-spectrum coronavirus antiviral drug discovery. Expert Opin Drug Discov. 2019;14(4):397-412. doi:10.1080/17460441.2019.1581171.

[11] Touret F, de Lamballerie X. Of chloroquine and COVID-19. Antiviral Res. 2020;177:104762. doi:10.1016/j.antiviral.2020.104762.

[12] Colson P, Rolain JM, Raoult D. Chloroquine for the 2019 novel coronavirus SARS-CoV-2. Int J Antimicrob Agents. 2020;55(3):105923. doi:10.1016/j.ijantimicag.2020.105923.

[13] An J, Minie M, Sasaki T, Woodward JJ, Elkon KB. Antimalarial Drugs as Immune Modulators: New Mechanisms for Old Drugs. Annu Rev Med. 2017;68:317-330. doi:10.1146/annurev-med-043015-123453

[14] Li X, Ma X. Acute respiratory failure in COVID-19: is it "typical" ARDS? Crit Care. 2020 May 6;24(1):198.

[15] He R, Lu Z, Zhang L, Fan T, Xiong R, Shen X, Feng H, Meng H, Lin W, Jiang W, Geng Q. The clinical course and its correlated immune status in COVID-19 pneumonia. J Clin Virol. 2020 Apr 12;127: 104361.

[16] Vashishtha S, Broderick G, Craddock TJ, Fletcher MA, Klimas NG. Inferring Broad Regulatory Biology from Time Course Data: Have We Reached an Upper Bound under Constraints Typical of In Vivo Studies? PLoS One. 2015 May 18;10(5):e0127364

[17] Sims AC, Tilton SC, Menachery VD, et al. Release of severe acute respiratory syndrome coronavirus nuclear import block enhances host transcription in human lung cells. J Virol. 2013;87(7):3885-3902. doi:10.1128/JVI.02520-12.

[18] Nikitin A, Egorov S, Daraselia N, et al: Pathway studio - The analysis and navigation of molecular networks. Bioinformatics 19:2155–2157, 2003. doi: 10.1093/bioinformatics/btg29

[19] Guimerà R, Sales-Pardo M. Missing and spurious interactions and the reconstruction of complex networks. Proc Natl Acad Sci U S A. 2009;106(52):22073-22078. doi:10.1073/pnas.0908366106

[20] Sedghamiz H, Morris M, Craddock TJA, Whitley D, Broderick G. Bio-ModelChecker: Using Bounded Constraint Satisfaction to Seamlessly Integrate Observed Behavior With Prior Knowledge of Biological Networks. Front Bioeng Biotechnol. 2019;7:48. Published 2019 Mar 26. doi:10.3389/fbioe.2019.00048.

[21] Chu, G., Garcia de la Banda, M., Mears, C. et al. Symmetries, almost symmetries, and lazy clause generation. Constraints 19, 434–462 (2014). https://doi.org/10.1007/s10601-014-9163-9

[22] Klarner H., Streck A., Šafránek D., Kolčák J., Siebert H. (2012) Parameter Identification and Model Ranking of Thomas Networks. In: Gilbert D., Heiner M. (eds) Computational Methods in Systems Biology. CMSB 2012. Lecture Notes in Computer Science, vol 7605. Springer, Berlin, Heidelberg.

[23] Nethercote N, Stuckey PJ, Becket R, Brand S, Duck GJ, Tack G. MiniZinc: Towards a standard CP modelling language. In C. Bessiere, editor, Proceedings of the 13th International Conference on Principles and Practice of Constraint Programming, volume 4741 of LNCS, pages 529–543. Springer, 2007.

[24] Chen, M. S., Wang, H. F., Hwang, C. P., Ho, T. Y., & Hung, C. H. (2016). A variational bayesian approach for unsupervised clustering. In Lecture Notes in Electrical Engineering (pp. 651-660). (Lecture Notes in Electrical Engineering; Vol. 375). Springer Verlag. https://doi.org/10.1007/978-981-10-0539-8_63.

[25] Subedi S, McNicholas PD. Variational Bayes approximations for clustering via mixtures of normal inverse Gaussian distributions. Adv Data Anal Classif 8, 167–193 (2014). https://doi.org/10.1007/s11634-014-0165-7.

[26] Caocci G, La Nasa G. Could ruxolitinib be effective in patients with COVID-19 infection at risk of acute respiratory distress syndrome (ARDS)? [published online ahead of print, 2020 May 14]. Ann Hematol. 2020;1-2. doi:10.1007/s00277-020-04067-6.

[27] Miller R, Englund K. Clinical presentation and course of COVID-19. Cleve Clin J Med. 2020 May 5. doi: 10.3949/ccjm.87a.ccc013. [Epub ahead of print]

[28] Gu J, Gong E, Zhang B, et al. Multiple organ infection and the pathogenesis of SARS. J Exp Med. 2005;202(3):415-424. doi:10.1084/jem.20050828

[29] Sims AC, Burkett SE, Yount B, Pickles RJ. SARS-CoV replication and pathogenesis in an in vitro model of the human conducting airway epithelium. Virus Res. 2008;133(1):33-44. doi:10.1016/j.virusres.2007.03.013.