Award Number: W81XWH-15-1-0583

TITLE: High-Fidelity Design of Multimodal Restorative Interventions in Gulf War Illness

PRINCIPAL INVESTIGATOR: Darrell Whitley, PhD

CONTRACTING ORGANIZATION:

Colorado State University
601 S. Howe Street
Fort Collins, Colorado, 80523-2002

REPORT DATE: December, 2019

TYPE OF REPORT: Final

PREPARED FOR: U.S. Army Medical Research and Materiel Command
Fort Detrick, Maryland 21702-5012

REPORT DOCUMENTATION PAGE

Form Approved
OMB No. 0704-0188

1. REPORT DATE *(DD-MM-YYYY)*	2. REPORT TYPE	3. DATES COVERED *(From - To)*
DECEMBER 2019	FINAL	30SEP2015 - 29SEP2019

4. TITLE AND SUBTITLE
High-Fidelity Design of Multimodal Restorative Interventions

in Gulf War Illness

5a. CONTRACT NUMBER

5b. GRANT NUMBER
W81XWH-15-1-0583

5c. PROGRAM ELEMENT NUMBER

6. AUTHOR(S)
Whitley, L.D. (with G Broderick)

5d. PROJECT NUMBER

5e. TASK NUMBER

5f. WORK UNIT NUMBER

7. PERFORMING ORGANIZATION NAME(S) AND ADDRESS(ES)

COLORADO STATE UNIVERSITY.
601 S. Howe Street
FORT COLLINS, CO 80523-2002

8. PERFORMING ORGANIZATION REPORT NUMBER

9. SPONSORING / MONITORING AGENCY NAME(S) AND ADDRESS(ES)
U.S. Army Medical Research
and Materiel Command
Fort Detrick, MD 21702-5012

10. SPONSOR/MONITOR'S ACRONYM(S)

11. SPONSOR/MONITOR'S REPORT NUMBER(S)

12. DISTRIBUTION / AVAILABILITY STATEMENT

Approved for public release; distribution unlimited.

13. SUPPLEMENTARY NOTES

14. ABSTRACT

The objective of this project is to refine models of immune and endocrine regulatory dysfunction developed under W81XWH-10-1-0774 by Dr. Broderick and Dr. Craddock.

We have made progress towards the overall objectives of the grant and laid a foundation for future work. Specially, we improved an optimization tool for constraint satisfaction that is used to mathematically optimize models of biological regulatory networks. We also worked with Dr. Broderick's team to apply these methods to specific biological networks. We had previously optimized new regulatory networks that can capture cyclic behavior in the Hypothalamic-Pituitary Adrenal (HPA) axis model. We have now developed models of other types of biological regulatory networks. Currently, the methods that the team has developed work only on discrete finite models. Our last major research task was to apply machine learning methods to model discovery for models using continuous parameter models.

15. SUBJECT TERMS
GWI; Hypothalamic-pituitary-adrenal (HPA) axis

16. SECURITY CLASSIFICATION OF:			17. LIMITATION OF ABSTRACT	18. NUMBER OF PAGES	19a. NAME OF RESPONSIBLE PERSON
a. REPORT U	b. ABSTRACT U	c. THIS PAGE U	UU	16	USAMRDC
					19b. TELEPHONE NUMBER *(include area code)*

Standard Form 298 (Rev. 8-98)
Prescribed by ANSI Std. Z39.18

Table of Contents

Introduction and Overview

Gulf War Illness (GWI) is a complex illness with symptoms presenting across several principal regulatory systems including immunologic and endocrine components. We are working on concepts developed by Dr. G. Broderick which propose that GWI might involve a chronic imbalance in co-regulation between the nervous, endocrine and immune systems [1, 2, 3, 4]. To explore this idea, we needed to develop regulatory models that function as dynamical systems. Having developed these dynamical system models, we needed to extract information about the computational behavior of these systems.

One major goal of the last year has been to continue to develop new modeling methods for parameterized dynamical systems models. The methods we have used during most of this project are based on optimization methods that are closely associated with logic based systems that use discrete logical operations to describe the behavior of biological regulatory networks. We also invested significant time in improving the description of the regulatory networks during the early phases of this research. This has allowed us to express constraints more precisely and thus to reduce the computation time needed to describe and construct larger biological network models. Working with Dr. Broderick's team, we also produced constraint-based software tools for the identification of discrete regulatory logic parameters from sparse and incomplete data (**Sedghamiz et al., 2019a, Sedghamiz et al., 2019b**)[18,19]. These models have become more efficient and practical over the life of this grant, and are now being used by Dr. Broderick and his team. Dr. Broderick's and Dr. Craddock's teams have also studied the departure of model predictions from experimental data. Furthermore, they have also looked at how these models can be used to evaluate drug combinations to assess how these might or might not be translated into actionable pharmacological treatments.

This last year we started looking at how we can use more complex and powerful modeling tools. We have worked with methods borrowed from the field of machine learning, and from neural networks in particular. These methods can be used to go beyond constraint based optimization methods for constructing predictive models. We know that the problem of fitting our predictive models to the experimental data is under constrained. By "under constrained," we mean that can find multiple models that would (equally) match the available data. Indeed, we have found that this is true. This happens in part because the data is incomplete and partial. With the constraint based methods we have been using we are mostly limited to asking if there is a feasible solution. Also, the constraint based methods can only yield discrete answers under discrete conditions. For example, we can only assert that a particular agent in our model is present, such as cortisol, or it might be measured and described with a few discrete options. (For example, we might describe cortisol as 1) high, 3) normal, 4) low.) With these new neural network methods we have developed, we can now find multiple models that fit a given set of partial data and we can utilize more precise numerical measurements. This has advantages and disadvantages, but developing such models was one of the goals of the original proposal. At the same time, we have found we can still impose discrete modeling requirements and this does not disrupt the neural networks ability to learn a model that fits the available data. Please consult Dr. Craddock and Dr. Broderick's report for the biological implications of this research. We will describe some of the ways we are developing ``neural network like" computational models that can also capture the computational behavior of regulatory networks.

Keywords: dynamical systems, regulatory immune system models, stable attractors, constraint propagation, neural networks.

Collaborating Organization: This work is being carried out in collaboration with Dr. G. Broderick under the associated grant number `W81XWH-15-1-0582.`

I. **Participants and Staffing**. In the final year of this work efforts were directed towards supporting model discovery and new constraint propagation methods.

II. **Personnel receiving support from this award**:

- Darrell Whitley, Faculty, Colorado State University
- Swetha Varadarajan, Ph.D. student, Colorado State University

  - Ms. Varadarajan was hired on January 1, 2018. She is a third year Ph.D. student in Computer Science.

Accomplishments:

Technological Developments

The task for Colorado State University was to support the development of optimization tools for the discovery and refinement of regulatory models being used and developed by Dr. Broderick's team. One major goal of this work was to understand the stable attractors of these systems. We also wanted to develop methods that could scale up so as to be able to extract information from regulatory models with hundreds of components. We will summarize the major results we have achieved.

Year 1, A Brief Summary: At the beginning of the project, Dr. Broderick was exploring the idea that information about the distribution of ``local minima" or minima energy states of a biological regulatory model could provide information about stable states in biological systems. This idea was explored in publications prior to the current research by Dr. Broderick and has been supported by results from other researchers (for example [12,13]). Normally, finding all of the local optima of an entire state space has exponential cost, and the discovery of all local optima is generally limited to models with less than 30 to 40 variables which can be enumerated. We were able to recast the problem of finding minima energy states as the exploration of a k-bounded optimization problem. This is possible because we use pre-existing information to determine variable interaction in the regulatory model. We can do this even when the biological model is underdetermined. We were then able to determine that certain hyperplanes of the search space could not contain a minima energy state. In this way, we were able to dramatically prune the search space and still find all of the minimal energy states of the system. While the methods we developed were highly effective, this static analysis of minima energy states proved to be inadequate to model the dynamics of biological regulatory models. After this, in years 2 and 3 we largely focused on constraint optimization methods to provide a better understanding of the more dynamical behavior of the model space.

Year 2: Dr. Broderick's team first pointed out that the search for regulatory models based on observable data could be posed as a constraint satisfaction problem (**Corblin et al 2012**)[7,8,9]. Thus, a constraint satisfaction solver was introduced into the research project. We contributed to this work by developing theoretical results to reduce the runtime of the solver and to provide the runtime guarantees in form of complexity results. We also have extensive experience with related complete Satisfiability (SAT) solvers such as MiniSat. In this project we have specifically been using Chuffed (**Ohrimenko et al. 2009**) [15] which is part of the MiniZinc software package. In a joint publication with Dr. Broderick's team we also explored the use of Google-OR-tools and the OptiMathSat solver for constraint satisfaction (**Sedghamiz 2019a**)[18]. Ultimately, different tools had different advantages on different classes of problems. This approach was also used in another publication looked at the discovery of minimal intervention sets in biological regulatory networks (**Sedghamiz 2019b**)[19]. Due to the use of constraint satisfaction solvers as an optimization tool, these results largely apply to systems that use discrete multi-value logic (for another example see [10].)

These models assume that particular interactions between agents (e.g. drugs, hormones, etc) are largely known but that we need to discover the strength of the excitatory and inhibitory signals that control the system. Constraint based search can be very costly, particularly since we need to find a globally optimal solution. One way that we have been able to restrict the cost of the search was to limit the search to a

very restricted set of possible interaction weights between agents. This approach has been successful, and all of our publications in the last two years of the project have made use of these constraint based methods [16,17,18,19]. One of the basic contributions that we made to the project was creating and refining problem formulations to yield optimization problems that could be solved more efficiently.

The regulatory models that have been presented to our team are closely related to the discrete modeling framework introduced by Thomas [20,21,22] and refined by other researchers [6,14]. We can think of nodes in a regulatory network model as representing basic biological variables, such as CRH from the pituitary gland, Adrenocorticotropin Hormone (ACTH), and cortisol (CORT) which binds to glucocorticoid receptors (R). These specific variables are generally are considered to be key elements of a simple regulatory model of the Hypothalamic-Pituitary-Adernal (HPA) axis [5]. This is captured in the following figure, where Stress is modeled as an external factor acting of the HPA axis (**Sedghamiz 2018**)[17,18].

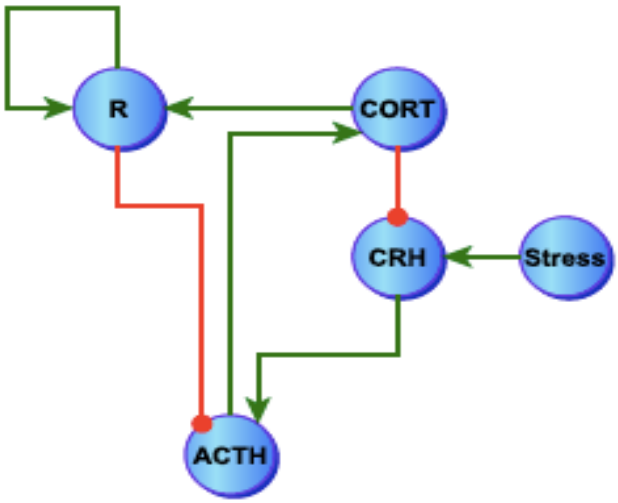

Figure 1: A simple regulatory model for the HPA axis.

As indicated in our previous report, when these systems are modeled using constraint based methods the variables in this model can only take on a discrete number of activation levels. For example, we might allow CORT to take on one of three levels (0 = low, 1 = nominal, 2 = high). This gives us a course grain representation that makes it easier to deal with variations in measurement, but it also makes all possible states of the total regulatory network both discrete and finite. Being discrete and finite is also necessary if we wish to apply constraint based optimization methods to the problem of modeling the dynamical behavior of the regulatory network. At the same time, it would be useful to also be able to have more continuous models of these regulatory networks. Methods borrowed from machine learning and neural networks give us the potential to describe and model these systems using continuous mathematics instead of discrete logical descriptions.

We can represent each node in the regulatory network by a variable X_i. The connections between variables are inhibitory or excitatory. The inhibitory and excitatory interactions are represented by a variable K_{ij} where variable *i* has an excitatory (positive) impact on variable *j* or an inhibitory (negative) impact on variable *j*. Variable X_i also has a threshold, such that variable X_i will only change state if there is sufficient inhibitory or excitatory signals from other connected variables.

The team at Colorado State University had developed an alternative method for modeling biological regulatory networks that utilizes continuous mathematics instead of discrete constraint satisfaction based methods. This work is motivated by the premise that the methods used to train neural networks in machine learning can be repurposed to also learn regulatory networks.

Neural Network Methods Adapted to Model Regulatory Networks.

Constraint based search can be very costly, particularly since we need to find a globally optimal solution. One way that we have been able to restrict the cost of the search was to limit the search to a very restricted set of possible interaction weights between agents. However, we have been developing new methods for learning this type of model that is much more efficient; by using methods from machine learning, we have the potential to find models very quickly and also to provide mechanisms to find multiple models that can account for the data. We will continue to emphasis however, that the problem is under-constrainted. Since the problem is under-constrainted, it is valuable to have methods that can generate multiple models that can equally account for the available data. We should also note that the following methods, at least for now, assume that updates in these biological models are synchronous [11].

In the first part of this section, we revisit some of the report from last year (2018) in order to set the stage for work carried out in 2019. Consider the following variation on the HPA axis model.

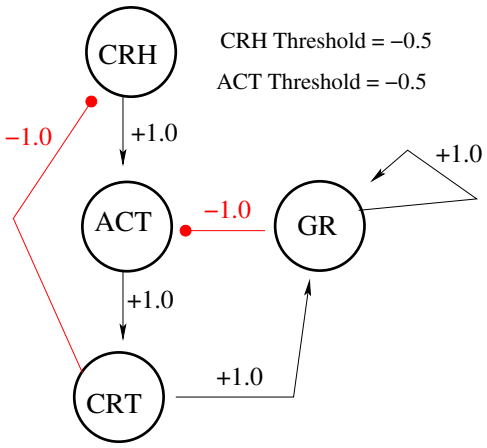

Figure 2: An alternative view of the simple regulatory model for the HPA axis.

This is the same as the previous HPA axis model (Figure 1) except the external **stress** has been removed and specific excitatory and inhibitory weights have been introduced in to the model. (In this model, CRT = CORT, ACT = ACTH, GR = R, and CRH =CRH). In order to convert this into a form that neural network methods can be applied, we first need to normalize the response thresholds (by making all thresholds 0) and reduce learning to the problem of discovering the weights between the agents (where agents are represented as the vertices). This can be done by introducing a vertex that is always "active." Such vertices are called "True Nodes" in the neural networks literature. For example, both CRH and ACT have a Threshold of -0.5 in this model. This means that CRH will become active even when no other agent acts on CRH; for example, CRH can "turn on" when CRT is inactive. By introducing a "True" vertex that is *always active,* thresholds are no longer required. If there exists a threshold denoted by T_d for vertex *d*, the threshold is transformed into a weight, $w_{t,d}$ where t denotes the "true node" and *d* denoted the vertex where the threshold is being replaced.

We next apply a method to convert the model into a vector-matrix form. The method we are using is a from of ``loop unrolling." We express the state of the agents (represented by the vertices) in a vector **v** where v_t denotes the activation levels of the vertices at time t. The activation levels at the next time step (time t+1) can also be expressed in vector form as $\mathbf{v_{t+1}}$. The weights on the edges connecting the vertices can now be expressed by a matrix M, such that

$$v_t M = \mathbf{v_{t+1}}$$

If we wish to add a nonlinear limiting (or "squashing") function to keep agent activation levels to a particular range, this can also easily be done: $S(v_t M) = \mathbf{v_{t+1}}$ where S is a limiting function (such as the logistics function) that asymptotically limits the upper and lower range of the activation states in vector **v.**

The idea of loop unrolling and of modeling the HPA network as a linear feedforward networks is captured in the following illustration:

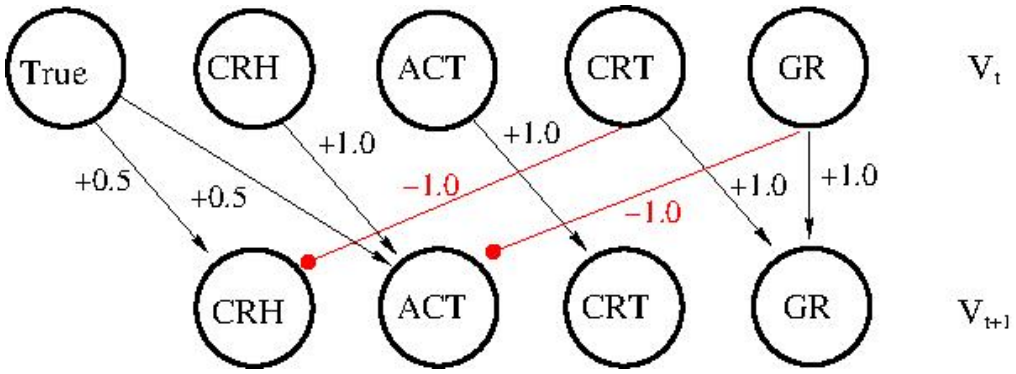

Figure 3: The alternative simple regulatory model for the HPA axis configured as a set of linear transitions.

As a linear system, the matrix **M** can be discovered or learned using the perceptron learning rule that is used to train linear neural networks. All we need to do this is to present examples of the transitions that are characteristic of the biological regulatory network. We should also note that the matrix **M** is sparse, just as we have indicated in all of the models presented here. However, we do not have to tell the system which connections are excitatory and which are inhibitory. Perception learning can recover this information from the training data. In the "unrolled" from, the HPA model also has exactly the same form as a single layer feedforward neural network.

We can use the following training data to train the model:

Input **v**	→	Output **v**
0001	→	1001
1001	→	1101
1101	→	1111
1111	→	0111
0111	→	0011
0011	→	0001

where the data is expressed in vector form and (for example) 0001 denotes: CRH = 0, ACT = 0, CRT = 0 and GR = 1. (Similarly 1101 denotes CRH = 1, ACT = 1, CRT = 0 and GR = 1.)

It should be stated clearly at this point that this new work breaks with our previous work. Do we assume that we already have some form of regulatory model, such as the model given in Figure 1, and that we will then search for the set of behaviors that can be supported by that model? Or do we assume that we first only have data and then search for a regulatory model that is consistent with the data. Probably neither approach is ideal. But to apply methods from machine learning we must assume that train data is available. Our assumption about data also raises other questions about how data should be collected, and even what is means to have "representative data." For example, with biological processes, different data sets might be collected at different time intervals resulting in different levels of granularity. Data collected at different

time intervals obviously results in regulatory models with different levels of granularity as well. These are issues we have frequently discussed with Dr. Broderick's team. And ultimately, these are issues for all models, regardless of whether they result from constraint satisfaction methods or machine learning. We do not resolve the question of granularity. The best we can do at this point is to accept some predetermined level of granularity. (Again, this granularity might also be *implicit* rather than *explicit* in the sense that the data itself encapsulates a particular level of granularity.)

With this training data we can recover the original model shown in Figure 1. But we also recover other models that have identical behaviors. Some are merely scalar transforms of the same model. But other models represent different functional forms. For example, in the training data used here, GR always has activation "1" and thus is under constrained. This model accounts for the available data just as well as any other model we have previously used to model the HPA axis and it does so with fewer interconnections. An example is given in Figure 3.

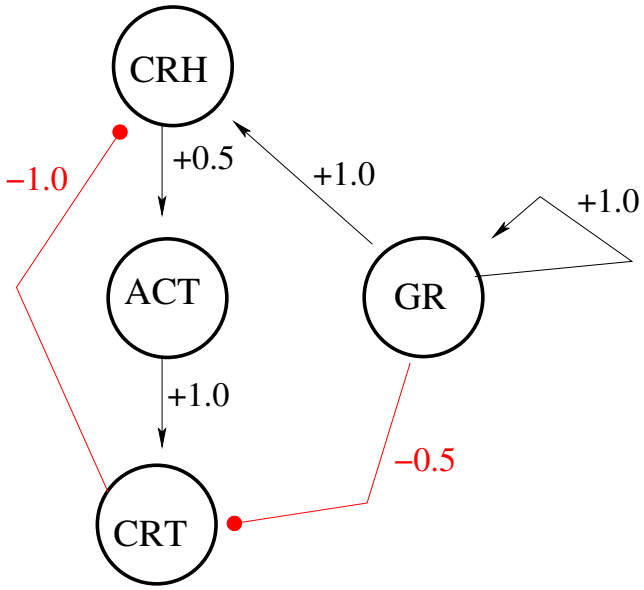

Figure 3: An alternative solution describing the simple regulatory model for the HPA axis.

The learning algorithm also discovered true simplifications of the model. For example, since GR is self excitatory and there are no inhibitory connections to GR in Figure 1, there is no need for CRT to excite GR in the HPA model and this can be removed with no impact on the computational behavior of the regulatory model, at least with regard to the mathematical model in Figure 1; the actual biology of course might be different. Only when additional agents are added to the HPA model might we also need the excitatory connection between CRT and GR. This represents a true simplification of the HPA model in is current form, and also highlights the need to consider the wider operating conditions of the regulatory networks.

Given sufficient training data, this representation of regulatory models can learn any linear regulatory system. However, the question of what constitutes *sufficient training data* is not trivial. We have noted on several occasions that given the amount of data that is available, modeling regulatory networks is under constrained. In the worse case, the size of a *complete set of training data* could be (virtually) exponential in size relative to the number of vertices in the regulatory model. In practice the problem of needing excessive amounts of training data seems unlikely however. We assume that regulatory networks have a limited number of behaviors that correspond to the specific behaviors we want to model.

Discrete vs Continuous Weights: Theoretically, the machine learning algorithm we are using is capable of learning arbitrary continuous weighted connections between agents (vertices). However, in the cases that we have studied, the model can also learn discrete weights. For example, assume we want all of the weights in the system to move with a step size of 0.1. This means that every weight in the final model will be some multiple of +/- 0.1. This causes the learning algorithm to find a finite set of weighted values that (in effect) form a "staircase" of what is effectively a multi-valued logic. In our experiments the number of weighted values is small. For example, a learning model might only use the set of weights drawn from the set: {-0.3, -0.2, -0.1, 0, 0.1, 0.2, 0.3}. This means that machine learning can also potentially also learn the same kind of restricted multi-valued logical models being learned using constraint satisfaction methods, but it can do so much faster, and it can automatically discover and adjust the level of discretization that is required. This discretization might also have some advantages in as much as it allows for qualitative descriptions of the connections (e.g., excitatory or inhibitory; low or nomimal or high). It also has the side effect of "covering" a large (indeed a staggeringly large) number of models which are all doing exactly the same computation with minor differences in the set of continuous weights between the variables.

The "CAMPFIRE" models: In this section, we report on new work that was designed to highlight the use of neural networks to learn models of abstract biology models given limited data. We stress these models of are artificial and idealized, but this is very intentional. We did not want to apply these methods to biological data at this point. Instead, this works highlights the potential of what can be done with adequate data, and need to collect better data for developing and evaluating biological systems models.

Again, the **CAMPFIRE** model is based on artificial data. **C, A, M**, and **P** are all agents corresponding to variables in the biological model. The particular variables **C, A, M**, and **P** are loosely based on CORT, ACTH, GR and CRH in the HPA axis. The variable **F** is always on and acts as "true node" in the neural model. The variables **I, R** and **E** tend to be semi-dependent variables driven by the **CAMP** variables, but these variables also impact the cycle time of the overall model.

To apply machine learning, we need training data. We assume we have the following observational data for the variables **v** = C, A, M, P, F, I, R, E. (Naming the variables, while quite informal, also makes it easy to track the interactions.)

Input **v**	→	Output **v**
00001000	→	10001010
10001010	→	11001011
11001011	→	11111001
11111001	→	01111100
01111100	→	00111001
00111001	→	00011000
00011000	→	00001000

Note that some of the **C,A,M,P** variable interactions mimic transitions found in the HPA model in Figure 1. Based on the data, the following learned matrix M represents a solution to the linear model

$$v_t M = \mathbf{v_{t+1}}$$

where the matrix M is given by:

	C	A	M	P	F	I	R	E
C	Z	1.2	Z	Z	Z	0.5	Z	Z
A	Z	Z	1.2	1.2	Z	Z	-1	Z
M	Z	Z	Z	1.2	Z	0.5	Z	Z
P	-1.1	Z	Z	Z	Z	Z	-1	Z
F	0.9	-0.1	0	0	0.9	-0.7	0.5	0
I	Z	Z	Z	Z	Z	Z	Z	1.6
R	Z	Z	Z	Z	Z	Z	Z	1.6
E	Z	Z	Z	Z	Z	Z	Z	Z

In this case, we assumed we knew which variables had potential interactions, but we did not include information as to which interactions were excitatory and which were inhibitory. This level of detail was learned. The character "Z" in the matrix indicates that these interactions were penalized to drive the interaction weight toward zero, but they were not fixed to be zero. This results in the following interaction model based on the data.

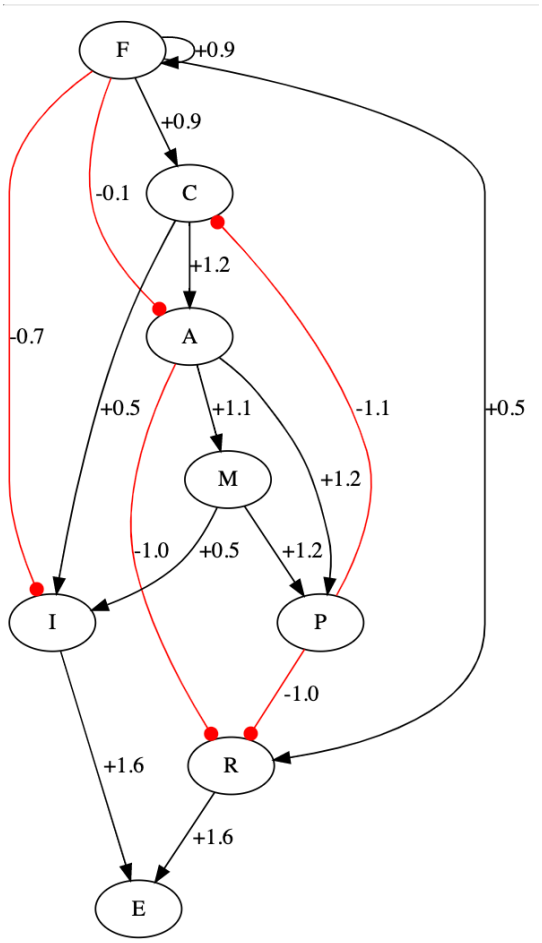

Figure 4: A solution yielding an interaction model for the CAMPFIRE data.

Having this model and the matrix allows us to expand the entire transition space for all combinations of variables. This is captured in the following network of state transitions.

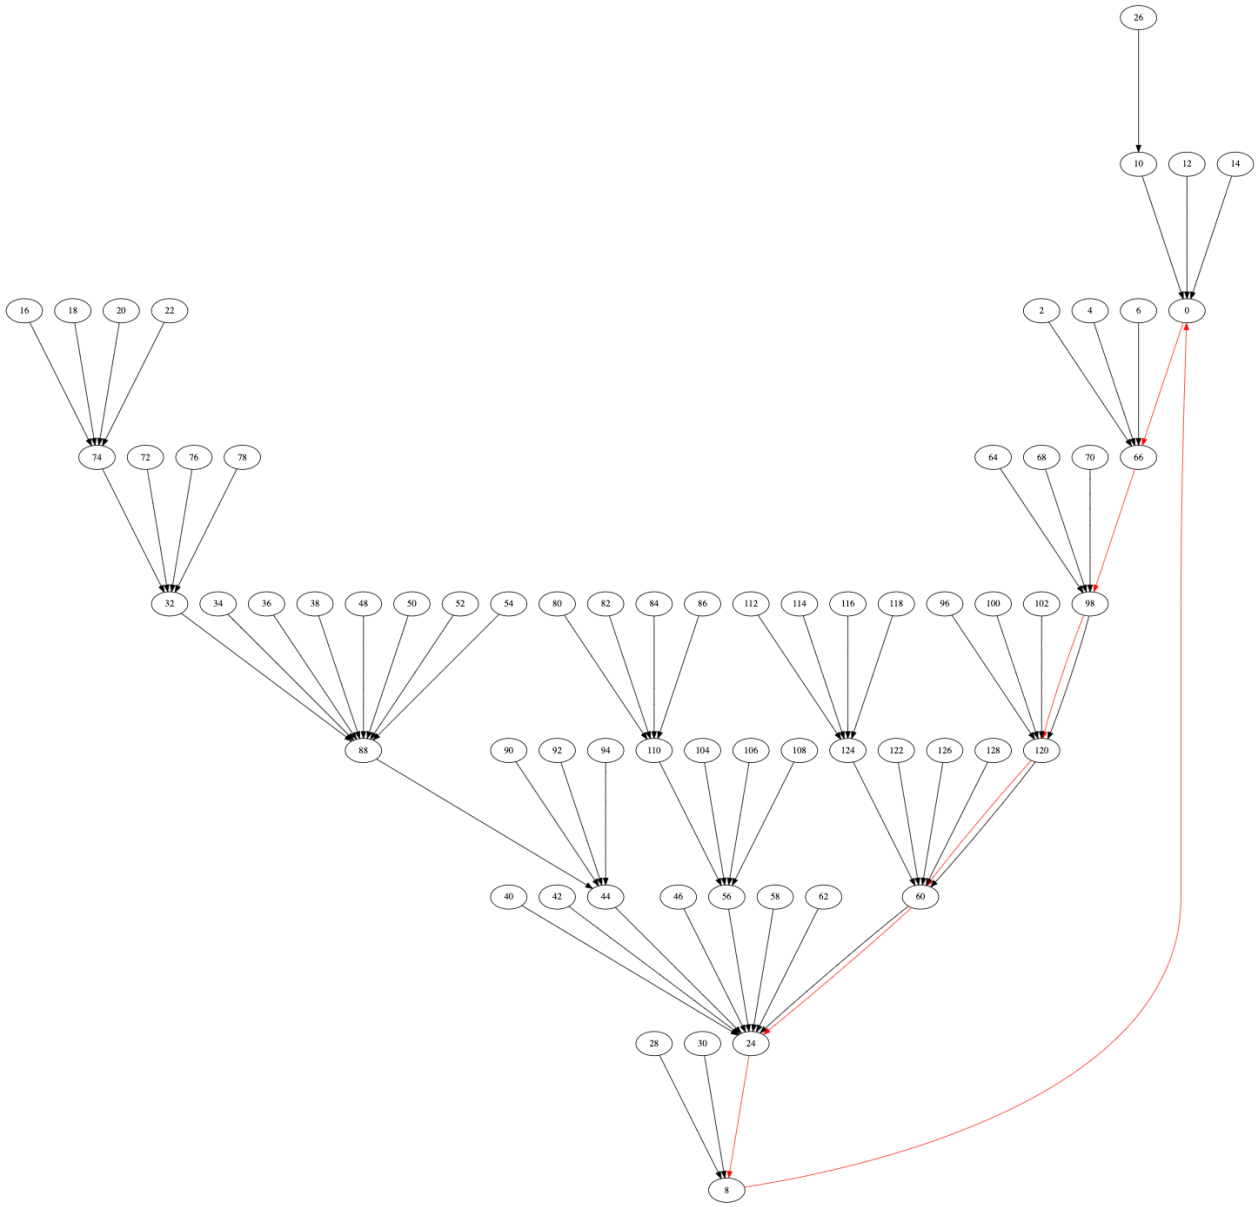

Figure 5: The full interaction graph for all of the transitions resulting from the static CAMPFIRE model given in Figure 4.

In this case, the vectors of bit strings over the CAMPFIRE variables are converted to integers. The red edges in this model connects the data used to train the network model, where that data include a cycle of repeating states included in the training data.

M**odel Flexibility:** It is possible to learn this model and similar models while making few or no assumptions about which variables interact. We have found it is enough to (slightly) penalize all variables to drive them toward zero if it is believed that there exists no interactions between specific variables. In this way, only the necessary variables interactions tend to be non-zero. But depending on the level of penalty, different models can be produced, all of which “fit” the training data. Different models can, of course, yield

very different networks of state transitions over all possible variable combinations. In other words, we can very easily obtain multiple models that fit the data. This is potentially useful. By making our assumptions stronger or weaker concerning which variable interact, we can generate alternative (biological) models very rapidly and also generating the entire network of state transitions for each model. This, in turn, allows us to generate alternative hypotheses that could potentially be tested by collecting new data that was not part of the training data.

Iterative Model Refinement: Biological researchers might prefer an approach where the biological model comes first, and then the behaviors of the model are captured using constraint satisfaction methods (which is what we have in effect done for the last two years). The alternative is to derive the model completely from data. However, the two approaches can be used in tandem. Given a proposed biological model and multi-value logic that describes the behavior of the model, we can use this model to generate data. We can then turn around and use machine learning to ask what other models display the same behaviors using that same data. This results in an iterative process that alternates between a model driven approach and a data driven approach to model refinement.

From Models to Logical Constraints: While working with these models, it became clear that the neural network was really learning a set of logical constraints. For example, for the CAMPFIRE model given in Figure 4, we can easily extract the following logical rules that also account for the training data:

SET OF LOGICAL CONSTRAINTS FOR THE "CAMPFIRE" DATA.

CONDITION	ACTION	
If (P=0) at time t	then (C = 1) at time t+1	else C = 0
If (C=1) at time t	then (A = 1) at time t+1	else A = 0
If (A=1) at time t	then (M = 1) at time t+1	else M = 0
If (A=1) or (M=1) at time t	then (P = 1) at time t+1	else P =0
If TRUE	then (F=1)	
If (C=1) and (M=1)	then (I=1) at time t+1	else I = 0
If (A=0) and (P=0)	then (R=1) at time t+1	else R = 0
If (R=1) or (I = 1)	then (E=1) at time t+1	else E = 1

Thus, the machine learning methods can provide us with a relatively simple way of not only discovering models from data, we can also convert those models into simple logical constraint-based rules of behavior. In many cases, these same kinds of rules can describe the weighted multi-valued models found by constraint satisfaction methods for the HPA axis and other biological regulatory networks that we have studied over the last three years.

This also provides a path to "explainable" models, which falls under the larger umbrella of "explainable Artificial Intelligence" and explainable machine learning. Nevertheless, we should stress that these particular logical constraints are relatively simple. However, the logical constraints are relatively simple only because the training data is binary, and thus only indicates that particular agents are active or not active (using a 1 or 0 representation). More work needs to be done to develop richer models capable of representing either continuous data inputs or much more detailed multi-valued logics. The machine learning methods we are using are fully capable to modeling continuous data inputs, but when using continuous data it will not be as simple to extract rules to describe model behaviors. Going forward, there may be advantages to using a combination of continuous data and multi-valued data with the goal of constructing more explainable models.

Impact:

- We contributed to the implement of the software tool, *Bio Model Checker (BioMC)* environment, now being used by non-programmers to develop high-fidelity models of biological signaling.
- The same *Bio Model Checker (BioMC)* environment has contributed to fault analysis techniques developed in the microelectronics industry to handle higher-resolution multi-valued logic required to adequately describe biological signaling at levels of biology other than the genome.
- Outside of biology, we have developed new optimization techniques that have the potential to greatly speed up constraint based optimization methods. So far, we have largely tested these methods on Maxiumum Satisfiably (the foundational problem in constraint satisfaction) but the methods generalize to other problems in constraint satisfaction.
- We have developed a neural network implementation for learning regulatory models. These models allow us to discover multiple regulatory models while making fewer assumptions about which variables interact. It also makes it unnecessary to make a prior distinctions between excitatory and inhibitory interactions, since these can be learned.
- We will also point to the impacts outline in the parallel report being submitted by Dr. Broderick and Dr. Craddock. While we cannot take direct credit for some of this work, we have supported the overall project over the life of the project.

What was the impact on other disciplines?

- Outside of biology, we have developed new optimization techniques that have the potential to greatly speed up constraint based optimization methods. So far, we have largely tested these methods on Maxiumum Satisfiably (the foundational problem in constraint satisfaction) but the methods also generalize to constraint satisfaction. Our original goal was to improve the constraint satisfaction solvers that we have been using in this research.

.

Changes/Problems: None.

Products

Publications:

Francisco Chicano, Gabriela Ochoa, Darrell Whitley, Renato Tinós (2019)
Quasi-Optimal Recombination Operator. European Conference on Evolutionary Computation in Combinatorial Optimization. 131-146, Springer. (Best Paper Award.)

Hooman Sedghamiz, Matthew Morris, Travis JA Craddock, Darrell Whitley, Gordon Broderick (2019)
Bio-ModelChecker: Using Bounded Constraint Satisfaction to Seamlessly Integrate Observed Behavior with Prior Knowledge of Biological Networks. *Frontiers in Bioengineering and Biotechnology* 7,48.

H. Sedghamiz, M. Morris, D. Whitley, TJA. Craddock, M. Pichichero, G. Broderick (2019).
Computation of robust minimal intervention sets in multi-valued biological regulatory networks. *Frontiers in Physiology* 10,241.

Hooman Sedghamiz, Matthew Morris, Travis JA Craddock, Darrell Whitley, Gordon Broderick. (2018) High-fidelity discrete modeling of the HPA axis: A study of regulatory plasticity in biology. *BMC systems biology* 12 (1), 76

W Chen, D Whitley, R Tinós, F Chicano (2018).
Tunneling between plateaus: improving on a state-of-the-art MAXSAT solver using partition crossover. Proceedings of the Genetic and Evolutionary Computation Conference, 921-928. (Best Paper Award, Genetic Algorithms Track.)

H. Sedghamiz, W. Chen, M. Rice, D. Whitley and G. Broderick (2017)
Selecting Optimal Models Based on Efficiency and Robustness in Multi-valued Biological Networks. IEEE 17th International Conference on Bioinformatics and Bioengineering (BIBE), 200--205.

Special Reporting Requirements: None

References.

1. Craddock TJ, Fritsch P, Rice MA Jr, del Rosario RM, Miller DB, Fletcher MA, Klimas NG, Broderick G. A role for homeostatic drive in the perpetuation of complex chronic illness: Gulf War Illness and chronic fatigue syndrome. PLoS One. 2014 Jan 8;9(1):e84839.
2. Craddock TJA, Harvey JM, Singh SJ, Rice MA Jr., del Rosario RM, Klimas NG, Fletcher MA, Broderick G. Fire in the Head: Exploring the Role of Homeostatic Drive in the Perpetuation of Neuroinflammation and Brain Disorders. Miami 2014 Winter Symposium: The Molecular Basis of Brain Disorders, Miami, FL, January 26-29, 2014. Poster session 2.
3. Fritsch P, Craddock TJ, del Rosario RM, Rice MA, Smylie A, Folcik VA, de Vries G, Fletcher MA, Klimas NG, Broderick G. Succumbing to the laws of attraction: Exploring the sometimes pathogenic versatility of discrete immune logic. Sys Biomed 2013, 1(3):1
4. Rice MA Jr, Craddock TJA, Folcik VA, del Rosario RM, Barnes ZM, Klimas NG, Fletcher MA, Zysman J, Broderick G. Gulf War Illness: Is there lasting damage to the endocrine-immune circuitry? Sys Biomed, 2014 2(4): 80-89.
5. Ben-Zvi A, Vernon SD, Broderick G. Model-based Therapeutic Correction of Hypothalamic Pituitary Adrenal Axis Dysfunction. PLoS Comput Biol 2009, 5(1): e1000273.
6. Bockmayr A, Klarner H, Siebert H. Time series dependent analysis of unparametrized Thomas networks. IEEE/ACM Transactions on Computational Biology and Bioinformatics, vol. 9, pp. 1338–1351, 2012.
7. Corblin F, Fanchon E, Trilling L, Chaouiya C, Thieffry D. Automatic Inference of Regulatory and Dynamical Properties from Incomplete Gene Interaction and Expression Data. Berlin, Heidelberg: Springer Berlin Heidelberg, 2012, pp. 25–30. [Online]. Available: http://dx.doi.org/10.1007/978-3-642-28792-3 4

8. Craddock TJ, Fritsch P, Rice MA Jr, del Rosario RM, Miller DB, Fletcher MA, Klimas NG, Broderick G. A role for homeostatic drive in the perpetuation of complex chronic illness: Gulf War Illness and chronic fatigue syndrome. PLoS One. 2014 Jan 8;9(1):e84839.

9. Craddock TJ, Del Rosario RR, Rice M, Zysman JP, Fletcher MA, Klimas NG, Broderick G. Achieving Remission in Gulf War Illness: A Simulation-Based Approach to Treatment Design. PLoS One. 2015 Jul 20;10(7):e0132774. doi: 10.1371/journal.pone.0132774.

10. Faur A, Naldi A, Chaouiya C, Thieffry D. Dynamical analysis of a generic boolean model for the control of the mammalian cell cycle," Bioinformatics 2006, 22(14): e124, [Online]. http://dx.doi.org/10.1093/bioinformatics/btl210

11. Garg A, Di Cara A, Xenarios I, Mendoza L, De Micheli G. Synchronous versus asynchronous modeling of gene regulatory networks, Bioinformatics 2008, **24(17)**: 1917, 2008. [Online]. Available: +http://dx.doi.org/10.1093/bioinformatics/btn336

12. Gupta S, Aslakson E, Gurbaxani BM, Vernon SD. Inclusion of the glucocorticoid receptor in a hypothalamic pituitary adrenal axis model reveals bistability. Theor Biol Med Model. 2007 Feb 14;4:8.

13. Kim LU, D'Orsogna MR, Chou T.Onset, timing, and exposure therapy of stress disorders: mechanistic insight from a mathematical model of oscillating neuroendocrine dynamics. Biol Direct. 2016 Mar 25;11(1):13.

14. 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. Lecture Notes in Computer Science, vol 7605. Springer, Berlin, Heidelberg

15. Ohrimenko O, Stuckey PJ, Codish M. Propagation via lazy clause generation, Constraints 2009, 14(3): 357–391. [Online]. Available: http://dx.doi.org/10.1007/s10601-008-9064-x .

16. Sedghamiz H, Chen W, Rice M Jr., Whitley D, Broderick G. Selecting Optimal Models Based on Efficiency and Robustness in Multi-valued Biological Networks. IEEE 17th International Conference on Bioinformatics and Bioengineering (BIBE), 2017; 00(0): 200-205.

17. Sedghamiz H, Morris M, Craddock TJA, Whitley D, Broderick G. High-fidelity discrete modeling of the HPA axis: A study of regulatory plasticity in biology. BMC Sys Biol. 2018; 12: 76.

18. 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. Frontiers in Bioengineering and Biotechnology, 2019a, Mar 26:7-48.

19. Sedghamiz H, Morris M, Whitley D, Craddock TJA, Pichichero M, Broderick G. Computation of Robust Minimal Intervention Sets in Multi-valued Biological Regulatory Networks. Frontiers in Physiol. 2019b. Mar 19;10:241.

20. Thomas R, Thieffry D, Kaufman M. Dynamical behaviour of biological regulatory networks--I. Biological role of feedback loops and practical use of the concept of the loop-characteristic state. Bull Math Biol. 1995 Mar;57(2): 247-276.

21. Thomas R, Kaufman M. Multistationarity, the basis of cell differentiation and memory. II. Logical analysis of regulatory networks in terms of feedback circuits. Chaos. 2001 Mar;11(1):180-195.

22. Thomas R. Regulatory Networks Seen as Asynchronous Automata: A Logical Description. J Theor Biol 1991; 153: 1-23.