

Image-based spatiotemporal causality inference for protein signaling networks

Xiongtao Ruan¹, Christoph Wülfing² and Robert F. Murphy^{1,3,4,*}

¹Computational Biology Department, School of Computer Science, Carnegie Mellon University, Pittsburgh, PA 15213, USA, ²School of Cellular and Molecular Medicine, University of Bristol, Bristol BS8 1TD, UK, ³Departments of Biological Sciences, Biomedical Engineering, and Machine Learning, Carnegie Mellon University, Pittsburgh, PA 15213, USA and ⁴Freiburg Institute for Advanced Studies and Faculty of Biology, Albert Ludwig University of Freiburg, Freiburg im Breisgau, Baden-Württemberg 79104, Germany

*To whom correspondence should be addressed.

Abstract

Motivation: Efforts to model how signaling and regulatory networks work in cells have largely either not considered spatial organization or have used compartmental models with minimal spatial resolution. Fluorescence microscopy provides the ability to monitor the spatiotemporal distribution of many molecules during signaling events, but as of yet no methods have been described for large scale image analysis to learn a complex protein regulatory network. Here we present and evaluate methods for identifying how changes in concentration in one cell region influence concentration of other proteins in other regions.

Results: Using 3D confocal microscope movies of GFP-tagged T cells undergoing costimulation, we learned models containing putative causal relationships among 12 proteins involved in T cell signaling. The models included both relationships consistent with current knowledge and novel predictions deserving further exploration. Further, when these models were applied to the initial frames of movies of T cells that had been only partially stimulated, they predicted the localization of proteins at later times with statistically significant accuracy. The methods, consisting of spatio-temporal alignment, automated region identification, and causal inference, are anticipated to be applicable to a number of biological systems.

Availability and implementation: The source code and data are available as a Reproducible Research Archive at http://murphylab.cbd.cmu.edu/software/2017_TcellCausalModels/

Contact: murphy@cmu.edu

1 Introduction

Identifying the signaling pathways that control cellular processes is a critical goal of biomedical research. Current approaches emphasize high-throughput experimentation and computational analysis to identify proteome-scale interaction networks and regulatory relationships. Among various computational analysis frameworks, causality inference has been recently applied to analysis of biological systems, (Chang *et al.*, 2015; Chindelevitch *et al.*, 2012; Welf and Danuser, 2014). There are two predominant characteristics of causal relationships: (i) Cause always precedes effect; (ii) Performing an action (at least one example is known to exist) on the cause would change the effect (Eichler, 2012). For temporal causal modeling, the Granger causality test (Granger, 1969) has been an important tool. The idea of Granger causality is that if we can learn a better linear model for one variable by including information about another variable than by not including that variable, then we

suppose there to be a causal relationship between them. Various variants and generalizations of Granger causality and applications for biological problems have been described (Fujita *et al.*, 2010; Shojaie and Michailidis, 2010). However, Granger causality is a pairwise test, which limits its application to build large scale networks where the interactions are much more complicated than pairwise relationships. Another popular framework for temporal causal inference is the dynamic Bayesian model, a framework for statistical machine learning. Dynamic Bayesian models can be used to treat the protein network as a probabilistic graphical model, and can learn the structure and parameters from the data; however, it is controversial to treat the relationships in probabilistic graphs as causal relationships (Pearl, 2009). When compared with the Granger causality test method, dynamic Bayesian models perform better for short time series (Zou and Feng, 2009). However, when the structure of the network is unknown, it is not trivial to learn it; this is especially true

when there are a lot of variables (e.g. proteins), as is usually the case for biological problems.

In biological systems, many researchers resort to the second characteristic of causality, that is, to intentionally change (or perturb) the amount of a gene or protein suspected to play a regulatory role (e.g. through mutation, gene knockout or knockdown) and observe the effect on the suspected target or behavior. This approach has been used for analysis of single targets (e.g. genetic screens, high throughput or high content drug screens) and for inferring static or dynamic regulatory networks from measurements of the expression of many genes under perturbed and unperturbed conditions (Segal et al., 2003). Although the perturbation approach has yielded many important results, the complexity of many cellular pathways has led to a growing sense that causal relationships may be better addressed through analysis of kinetics and heterogeneity without explicit perturbation (Welf and Danuser, 2014). In this case, the various learning frameworks mentioned earlier could be applied.

Fluorescence microscopy is particularly suited for seeking causal relationships without relying on perturbation (of course, labeling with fluorescent probes may itself cause a different type of perturbation, but this is constant over the course of imaging experiments). Fluorescent tags can be used to identify particular molecules (which may represent causes or effects) and follow them in time and space in many living cells. Fluctuation analysis can then be used to look for correlations in the dynamics of two events (Vilela et al., 2013). For example, Alon and colleagues (Farkash-Amar et al., 2014) identified specific proteins whose expression level and/or localization is correlated with differences in cell motility by taking advantage of natural variation. Danuser and colleagues (Machacek, et al., 2009) have used fluctuation analysis to learn the relationship between myosin II localization and various aspects of cell shape and motility. These approaches have largely focused on identifying individual factors that influence one particular process, rather than constructing a model that identifies key steps that influence different proteins at various times and places. They have also required the ability to image both the putative cause (e.g. tagged protein) and the putative effect (e.g. cell shape) in the same cell; this permits the identification of correlated fluctuations.

Here we present machine learning pipelines for inferring putative causal relationships in regulatory or signaling networks. We note a number of differences from previous related work. First, our method does not require that pairs of proteins be observed in the same cells; relaxing this requirement is very useful when increasingly large numbers of proteins are addressed in the same system. Second, the method does not require manual specification of regions of the cell in which effects may occur; it allows for unbiased discovery. Last, it yields a graph of potentially interconnected influences rather than a list of individual factors.

Building on recent work (Roybal et al., 2016), we applied these causal inference methods to study the dynamics of proteins involved in T cell signaling. This is an interesting and important system for studying a complex spatiotemporal process. T cells become activated through interaction with antigen presenting cells. This is primarily accomplished by the T cell receptor recognizing antigen-derived peptide presented on the antigen-presenting cell (APC) surface. However, engagement of costimulatory receptors that interact with ligand on the APC is required for efficient T cell activation. The most potent costimulatory receptor is CD28, which binds to the B7 family ligands (CD80 and CD86). T cell activation produces rapid and transient accumulation of actin at the T cell/APC interface (also called an immunological synapse). Since T cells express many signaling proteins that can interact in complex manners, how the numerous individually well-characterized proteins integrate into a coherent system

of signal transduction is not well understood. It therefore represents a good test system for seeking causal relationships. We constructed models learned from movies of cells exposed to full stimulation and found a number of interesting relationships consistent with prior knowledge. Furthermore, applying the model to data for cells in which stimulation was partially blocked revealed that the model could make predictions about protein dynamics with statistically significant accuracy.

During initial analysis of the T cell movies, we found that several proteins displayed quite similar spatiotemporal patterns; this is probably due to the binding of these with each other and formation of persistent complexes. The spatiotemporal patterns of these proteins are therefore highly correlated, and thus the pattern of one can easily be predicted from the pattern of another. This is a common complication in causal analysis: the inability to distinguish correlation from causation. To deal with this situation, we eliminated correlated proteins using clustering when constructing our models.

It is important to note that the causal relationships mentioned below are the result of statistical inference and are intended to be read as putative causal relationships even when this is not explicitly stated.

2 Materials and methods

2.1 Data

We used the data and methods previously described (Roybal et al., 2016) for creating maps of the spatial distribution of ten different sensors under two conditions. The maps consist of the probability that a molecule of a given protein would be found in each of 6628 voxels in a standardized half-ellipsoid template; the flat surface of the template corresponds to the synapse. The values thus reflect relative, not absolute, concentrations. Maps were originally created for twelve time points relative to synapse formation, but we used only the first 10: -40, -20, 0, 20, 40, 60, 80, 100, 120 and 180 s for causal inference. We created maps for two new proteins: Ezrin and PIP2, using the same methods, and also slightly improved the alignment algorithm. Both the original and refined maps are available from <http://murphylab.web.cmu.edu/data/TcellModels>.

2.2 Region definition

Each voxel in the template was represented by a vector of length 120 containing the amounts of each of 12 sensors at ten time points. These vectors were grouped into voxel types by k -means clustering. Each connected component (using 26-connectedness) with the same cluster type and containing more than four voxels was defined as a region.

2.3 Causal graphical process model

2.3.1 Graph representation

A graph was created with a node for each region for each sensor; the node values are the concentration of all voxels in the region for that sensor. We consider that observing the graph yields a set of node values

$$\mathbf{x} = (x_1, \dots, x_n)$$

where n is the number of nodes.

For observation of a time series graph, we represent the nodes at time point k as

$$\mathbf{x}[k] = \left(x_1^{(k)}, \dots, x_n^{(k)} \right)$$

2.3.2 Causal graph process model

Given these values, we wish to construct an adjacency matrix A that captures the relationships between nodes, as well as to find a

function that describes how the node values change over time. As one approach, we chose to use Causal Graph Process (CGP) modeling (Mei and Moura, 2017) to do this. In the CGP, the function is defined as a polynomial in \mathbf{A} , and the order depends on the lag M :

$$\mathbf{x}[t] = w[t] + \sum_{i=1}^M P_i(\mathbf{A})\mathbf{x}[t-i] = w[t] + (c_{10}\mathbf{I} + c_{11}\mathbf{A})\mathbf{x}[t-1] \\ + (c_{20}\mathbf{I} + c_{21}\mathbf{A} + c_{22}\mathbf{A}^2)\mathbf{x}[t-2] + (c_{M0}\mathbf{I} + c_{M1}\mathbf{A} + \dots + c_{MM}\mathbf{A}^M)\mathbf{x}[t-M]$$

where $w[t]$ is the residual for $\mathbf{x}[t]$ (the amount of $\mathbf{x}[t]$ not accounted for by the model), and c_{ij} is the coefficient term for the j th order for the polynomial $P_i(\mathbf{A})$. We can write c_{ij} more compactly as

$$\mathbf{c} = (c_{10}, c_{11}, c_{20}, c_{21}, c_{22}, \dots, c_{MM})$$

Given this formulation, the goal is to find \mathbf{A} and c_{ij} ($i = 1, \dots, M, j = 0, \dots, M$). This is done by solving the optimization problem

$$\min_{\mathbf{A}, \mathbf{c}} \frac{1}{2} \sum_{k=M+1}^K \left\| \mathbf{x}[k] - \sum_{i=1}^M P_i(\mathbf{A})\mathbf{x}[k-i] \right\|_2^2 \\ + \lambda_1 \|\text{vec}(\mathbf{A})\|_1 + \lambda_2 \|\mathbf{c}\|_1$$

Where λ_1 and λ_2 are regularization parameters. $\text{vec}(\mathbf{A})$ means reshaping matrix \mathbf{A} as a vector, and K is the number of time points.

The problem is non-convex. As described previously in Mei and Moura (2017), we first solve $P_i(\mathbf{A})$, then estimate \mathbf{A} from $P_i(\mathbf{A})$, and then estimate \mathbf{c} from $P_i(\mathbf{A})$ and \mathbf{A} . Projected gradient descent (Figueiredo et al., 2007) is applied for each step. In order to prevent self-prediction, we add a constraint that the block of connections between nodes of the same protein is set to zero. This is easy to implement by setting these values to zeros in each iteration.

2.4 Constrained elastic net regression

2.4.1 Granger causality test

Pairwise Granger causality tests were performed for each pair of nodes. After calculation of the P -values, Bonferroni-Holm correction for multiple hypothesis tests was performed. A threshold α on the corrected P -values was used to identify significant pairs for constraining below. In this step, we also considered at which lag time the prediction is strongest by comparing the cross correlation with different lag times. The sign of the relationship (slope) from the linear models was also recorded to form the constraints in the next step.

2.4.2 Constrained elastic net regression

For identified pairs, if they have a positive linear relationship, then we forced the element in the adjacency matrix to be greater than b ($b > 0$), otherwise it is forced to be $< -b$. The basic optimization problem is

$$\min_{\mathbf{A}_1, \dots, \mathbf{A}_M} \frac{1}{2} \sum_{k=M+1}^K \left\| \mathbf{x}[k] - \sum_{i=1}^M \mathbf{A}_i \mathbf{x}[k-i] \right\|_2^2 \\ + \lambda_1 \sum_{i=1}^M \left\| \text{vec}(\mathbf{A}_i) \right\|_1 + \lambda_2 \sum_{i=1}^M \left\| \mathbf{A}_i \right\|_F^2$$

Subject to

$$\mathbf{A}_i(k, l) \geq b, \mathbf{A}_j(m, n) \leq -b$$

where nodes k, l are identified positive causal relationships at a particular lag i , and nodes m, n are identified negative causal relationships at

lag j , and λ_1 and λ_2 are regularization parameters. $\|\cdot\|_F^2$ means the Frobenius norm of a matrix.

We used the alternating direction method of multipliers algorithm (Boyd et al., 2011) to determine the adjacency matrices. One difference from CGP is that constrained elastic net regression (CENR) learns separate adjacency matrices for each time lag rather than a single adjacency matrix. Here we prevented self-prediction by setting the diagonal values to zeros after each iteration.

2.5 Grouping proteins

To avoid effects of correlations in spatiotemporal patterns that would interfere with the causal inference process, proteins were grouped based on hierarchical clustering. To do this, each protein was represented by the amount of that protein in each cell region at all-time points. Three similarity measures were used to construct dendrograms and decide the number of clusters: city block distance, Euclidean distance, and correlation coefficient. The mean kinetics of all proteins in a cluster were used to represent that cluster for causal model construction (for display, the cluster was given the name of the alphabetically first protein in it).

2.6 Parameter estimation

The best free parameter combinations for each method (λ_1 and λ_2 for CGP, α, b, λ_1 and λ_2 for CENR) were chosen using 5-fold cross validation over the kinetics for single cells. For construction of each training and test set, the mean kinetics of each region for each protein were calculated. The parameter combination giving the lowest average root mean normalized squared error (RMSE) over all proteins was chosen.

2.7 Timing of important events

To aid in interpretation of the model, we sought to identify specific time points at which effects captured by causal models were maximally observed. To do this (for lag M), we calculated for each time t the cross correlation between the concentrations of cause X for time points $t-M, \dots, t$ and the concentrations of effect Y for time points $t, \dots, t+M$. If the value of \mathbf{A} was positive for this effect, we found the t that maximized this cross correlation, and when the value of \mathbf{A} was negative we found the t that minimized it (largest negative value).

After defining the time point t of important event, the trend for cause and effect was defined as the slope of the linear model of cause X for time points $t-M, \dots, t$ and of effect Y for time points $t, \dots, t+M$, respectively.

2.8 Permutation tests

To examine the generalizability of the models, we applied them to predict kinetics for the B7 blockade condition not used for training the models. In the test, in each iteration, we permute on the order of kinetics for each protein and each region (on Full stimulus condition) in order to disrupt the association between different proteins and regions. Then We used the node values from the B7 blockade maps for time points $t-1, \dots, t-M$ to predict values for t ($t = M+1, \dots, K$) using the models trained with the full stimulus maps. As a measure of error of these predictions, the normalized RMSEs were computed between the predicted time series and the actual one, for all three regions for each sensor. To estimate the significance of obtaining a given RMSE for a given protein, permutation tests were performed in which predictions for a given protein in each of the three regions were randomly drawn from the values at all-time points in all regions for that protein. This process was repeated




Fig. 1. Region definitions. Regions within the standardized template found by voxel clustering and connected component analysis; different colors represent different regions, and magenta outside the cell indicates the background

50 000 times to obtain a background distribution of the prediction errors. The P -values for the predictions were corrected using the Bonferroni-Holm method and the corrected P -values below 0.05 were considered significant.

3 Results

3.1 Creating graphs of protein spatiotemporal interactions

3.1.1 Defining regions

We begin with maps of the relative concentration of twelve different proteins in each voxel of a standardized T cell template at various times before and after synapse formation (Roybal et al., 2016). These maps were generated from many movies of many sets of cells, each set expressing (at near the endogenous level) a different fluorescently tagged protein. Each movie showed an individual T cell forming a synapse with an antigen-presenting cell. The position and time of synapse formation was identified, the T cell boundary was found in each frame, and the fluorescence distribution was morphed to a standardized template. The amount of protein in each voxel of the standardized template was then averaged across many cells to provide a ‘map’ of the concentration.

If we consider the amount of each sensor in each position of the cell to form a fully connected graph, we can consider how the sensors change over time to be dictated by some process operating on that graph. This is the framework of the two methods: CGP and CENR method. However, if we treat each voxel separately, that is, as a node in the graph, then there will be 79 536 nodes (6628 voxels \times 12 proteins); learning the model would require estimating an adjacency matrix of over 6 billion values for CGP and 12 billion for CENR. Considering that nearby voxels have similar spatiotemporal patterns in the same process, we therefore chose to consider a smaller number of multi-voxel regions and defined these as set of connected voxels that behaved similarly: that is, those voxels that had similar kinetics of all sensors at all-time points. Note that this does not mean that all sensors were equal or that the amounts do not change, but rather only that if a sensor changes, it changes similarly among all voxels in a region. Possible regions were found using clustering of voxels into various numbers of clusters (see Section 2). Three clusters were observed to give the highest score by the Calinski Harabasz criterion, and therefore three voxel types (which happened to yield three regions) were used; the regions are shown in Figure 1.

3.1.2 Hierarchical clustering

As discussed earlier, we used clustering to eliminate spurious causal relationships due to high correlations among protein patterns. However, different similarity measures yielded different apparent numbers of clusters. For example, using city block distance, the clustering results shown in Figure 2 were obtained. Based on the figure, we created one group consisting of ARP3 and HS1 and a second group consisting of CP α 1, Coronin1A, WASP and WAVE2. As discussed in Section 2, these groups were represented by the average across all proteins in the group.




Fig. 2. Hierarchical clustering using city block metric. **Left:** The dendrogram shows the clustering of the proteins using the city block metric as described in the text. **Right:** the image shows the pairwise distance between each pair of proteins. The cophenetic coefficient of the dendrogram is 0.8866




Fig. 3. Adjacency matrix from the CGP. The X axis shows the cause and the Y axis shows the effect. The color of each element represents the estimated strength of the cause and effect relationship between a pair. ARP3 and CP α 1 each represent groups with multiple proteins

3.1.3 Building models

We next considered what temporal resolution we should attempt to capture within the graph process. Since changes in actin and its regulators happen rapidly (on the time frame of seconds to tens of seconds), and since we used only ten time points, we chose to learn models that for each time point depended on only the two prior time points ($M = 2$).

CGP model

Given the kinetics for each protein or group, we constructed CGP models and optimized parameters as discussed in Section 2. The resulting model consists of an adjacency matrix A , which provides the overall estimate of the effect of each node on each other, and the coefficients of the polynomials P that operate on it. The adjacency matrix is shown in Figure 3. The data (through the parameter optimization process) do not yield a very sparse model, suggesting there may be a complicated network among the proteins.

CENR model

Based on the threshold on corrected P -values, we identified 8 pairs of strong potential causal relationships. Based on the constraints, we trained the model with the regularization parameters. Again, we also suppressed self-prediction by setting the diagonal elements to zero. For this model, we learned one adjacency matrices for each lag time. These are shown in Figure 4. Despite our attempt to focus the model on strong effects (and similar to the case for the CGP model), the adjacency matrices are not very sparse.

3.2 Major findings

3.2.1 Major causal relationships

First, by hierarchical clustering, we identified proteins that have similar kinetics across all regions, suggesting that they are largely present in persistent complexes. There is prior evidence to suggest that this may be the case. *Urano et al. (2003)* has shown that HS1 is the upstream regulator for the activation of APR3, and activated ARP3 will bind to Actin and stabilize the nuclearization of Actin.

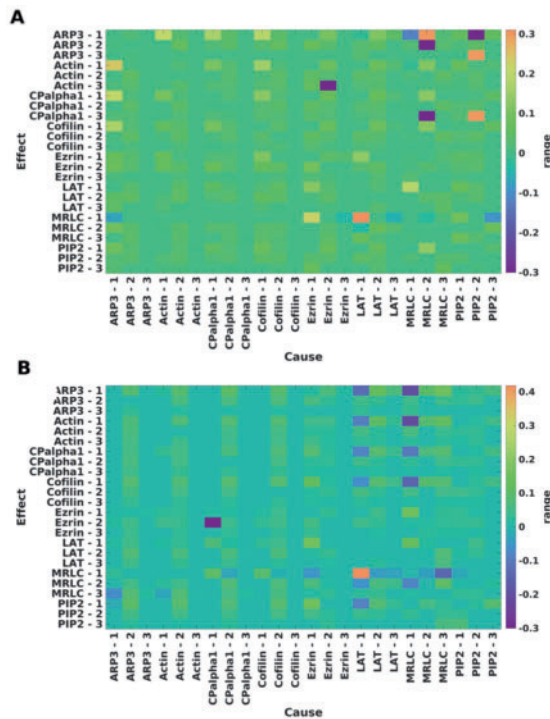


Fig. 4. Adjacency matrices from CENR. (A) lag time -1 . (B) lag time -2 . For both panels, the X axis shows the cause and the Y axis shows the effect. The color of each element in the matrix represents the strength of the cause and effect relationship between a pair at that lag time. ARP3 and CP α 1 each represent groups with multiple proteins




Fig. 5. Example kinetics for putative cause and effect relationships identified by CENR. The four pairs with the largest absolute values in the adjacency matrix are shown. In each panel, the x-axis shows time relative to synapse formation (seconds) and y-axis shows relative concentration. The cause protein and region are shown above each panel, followed by 'Pos' or 'Neg' to indicate the direction of the regulation and then the effect protein and region. The kinetics for the cause proteins are shown in solid black lines and the kinetics for the effect are shown in dotted red lines. Note the lag between changes in the cause and changes in the effect

WASP and WAVE2 are also upstream regulators of ARP3 (*Burkhardt et al., 2008*). Coronin1A is known as an Actin-binding protein that regulates actin through the coordination with ARP2/3 complex and cofilin (*Gandhi et al., 2009*). Therefore, our clustering results are consistent with prior studies. A comparison of the kinetics of cause and effect are shown for some of the largest relationships in Figure 5.

3.2.2 Region 1 as the most important region

In the adjacency matrices, we can see that the strongest relationships are mostly between proteins in region 1, the synapse region. This is the major region into which actin and its regulators are recruited and activated. This makes sense: this region is relatively small compare to other regions, and most proteins are accumulated in certain times, i.e. 0–40 s, i.e. the variation of the proteins is largest among the three regions. So clearly, these interesting patterns are consistent with what is known about actin dynamics in T cell synapse formation. However, we can also find some other pairs that are not in region 1, MRLC in the cortex (region 3) can predict the ARP3 group in both the synapse and the cortex. The dominant pair is that enrichment of myosin at the cell cortex drives accumulation of the core actin turnover machinery at the cellular interface. This suggests mechanical coupling between the cell body and the interface, an intriguing hypothesis that should and can be tested.

3.2.3 Critical time points

We can further explore these effects by identifying at which time the effect is maximally seen. This was done by finding, for positive and negative effects, respectively, the time point that had the largest or smallest (most negative) value in the cross-correlation function of the cause and effect proteins/regions. Using these, we can begin to link events into a network by examining whether the protein and regions that are the effect of an event at one time are the cause of an event at a later time. This is shown in Figure 6 for a larger number of relationships (a lower threshold on the adjacency matrix).

Inspecting this figure yields a number of interesting confirmatory and speculative conclusions. There are two major stages. First, most of the relations to actin in the synapse are observed to be maximal at the earliest two time points, i.e. an increase in cofilin and the ARP3 group in the synapse predicts the increase of actin and CP α 1 in the synapse. This is consistent with the rapid spreading of interface actin at these time points when the actin regulators would be expected to be most dynamic. Second, in the later stage, there is a reverse causal relationship for actin and its regulators such that a decrease of actin in synapse predicts a decrease of cofilin and the ARP3 group there. This is again consistent with the disassembly of actin after synapse formation. Moreover, proteins involved in immediate actin turnover, i.e. actin nucleation, capping and severing, are well interconnected yet mostly separated from three further

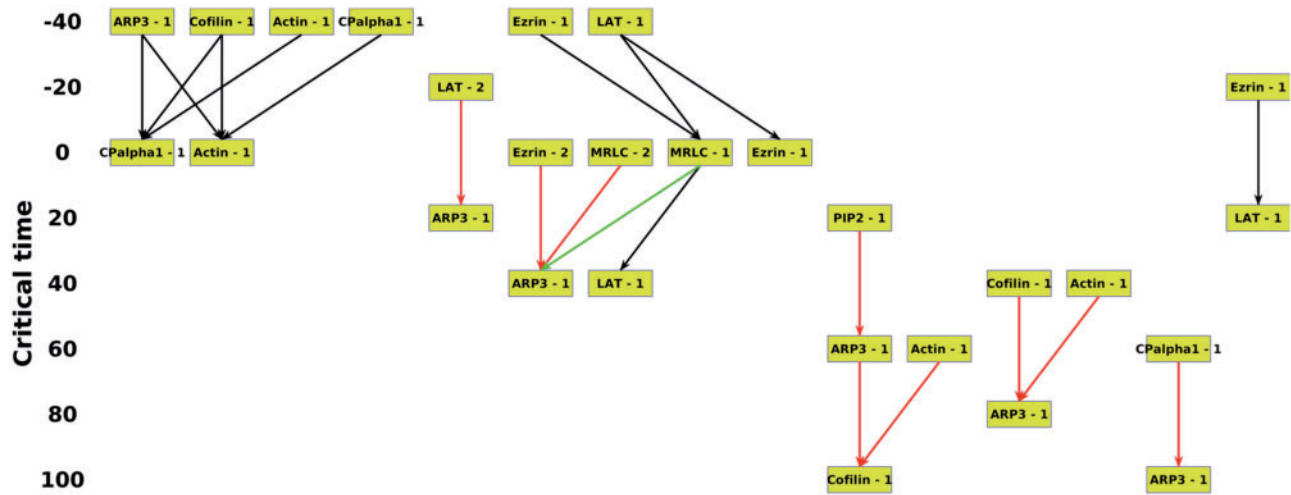


Fig. 6. Critical time points identified by the CGP. The time on the left when a putative causal relationship is maximally observed between a pair of proteins. Arrows indicate events and are colored by the sign of the change in the cause and effect: black correspond to both cause and effect increasing, red to both decreasing, green to cause increasing and effect decreasing, and blue to cause decreasing and effect increasing

upstream signaling regulators included in this analysis, LAT, Ezrin and myosin. This is to be expected. Surprisingly, the analysis identifies myosin as the key causal connector between upstream signaling and actin turnover. Although this observation needs to be confirmed by inclusion of a larger number of upstream signaling intermediates, a central role of myosin in coupling signaling to actin turnover is intriguing and testable.

3.2.4 Common pairs

To see how robust the relationships are across different methods, we compared the networks from the two methods (Fig. 7A). It is interesting that most relationships are found by both methods, especially the ones that are closely related to actin, such as those with the ARP3 group and cofilin. This suggests that these relationships are dominant in the actin dynamics in synapse formation, which is consistent with what we already know about actin dynamics. And on the other hand, it suggests both methods are reliable at finding the major relationships.

3.3 Performance of methods

As further evaluation of the ability of the methods to find potential causal relationships, we evaluated how well they found known relationships using Receiver-Operating Curve analysis. Based on database and literature searches, we identified potential relationships between these 12 proteins; however, since these sources often do not specify the directionality of relationships, we expressed them as an undirected graph (Fig. 7A). Since the strongest relationships occur in region 1, we considered whether those relationships could be found in region 1. The AUC values for recovery of the known relationships are shown in Table 1. Interestingly, both methods perform quite well; the AUC values of 0.709 and 0.644 (Case 1) are both much better than random.

3.4 Prediction for B7 data

The studies (Roybal et al., 2016) upon which this work is based also examined the effect of partial inhibition of T cell signaling upon sensor distributions. This inhibition was accomplished by blocking a secondary stimulation pathway, costimulation by CD28, and is referred to as B7 blockade. Maps of the same sensors were also generated under this condition. To explore how well the models learned

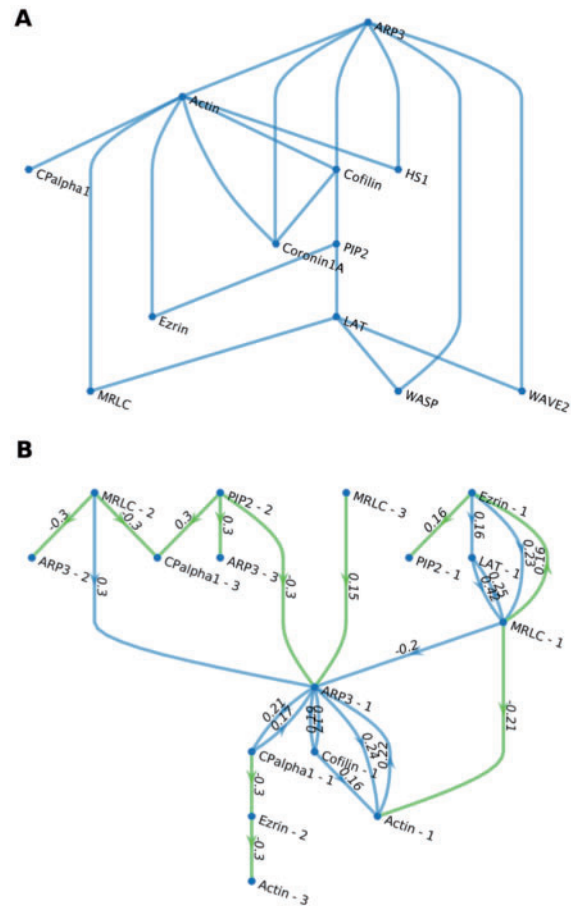


Fig. 7. Network comparison. (A) (partial) Ground truth, (B) Combined network. In (B), the edges are colored by which method identified a relationship: blue indicates the edge occurs in both two methods, green indicates the edge only was found by CENR. The numbers on the edges are the maximum absolute weight between the two methods

from the full stimulus conditions would be predictive of other conditions, we used the node values from the B7 blockade maps for time points $t - 2$, and $t - 1$ to predict values for t . Comparison of the actual and predicted kinetics are shown Figure 8. We then computed

Table 1. Comparison of different methods for different groups of proteins

	Ungrouped		Case 1		Case 2 (7)		Case 3 (25)	
	CGP	CENR	CGP	CENR	CGP	CENR	CGP	CENR
CV error	0.092	0.062	0.086	0.069	0.096	0.074	0.107	0.070
B7 error	0.119	0.080	0.120	0.085	0.117	0.081	0.123	0.082
AUC	0.629	0.565	0.709	0.644	0.589	0.580	0.614	0.589

Results are shown for the two methods for four different groupings of the seven similar proteins: ungrouped; (CP α 1, WASP, Coronin1A, WAVE2), actin, (ARP3, HS1); all grouped; and (CP α 1, WASP, Coronin1A, WAVE2, actin), (ARP3, HS1). CV error shows the mean testing error in 5-fold cross validation for the best parameter combination. The B7 prediction error was calculated after averaging any grouped proteins. The AUC values are for recovery of previously documented relationships in region 1.

the normalized mean squared errors between the predicted time series and the actual one, for each sensor, as a measure of error of the predictions. To test if the errors of these predictions were significantly less than expected at random, we performed permutation tests as described in Section 2. All predictions passed this test at the $P = 0.05$ level, indicating that at least the relationships between sensors and regions are conserved between the full stimulus and the B7 blockade conditions. The results suggest that the method can predict the turnover of actin and its regulators in the critical region even after perturbation.

3.5 Comparison of performance of different methods

Finally, we repeated the ROC curve analysis and the B7 prediction analysis for models constructed by both CGP and CENR either using all 12 proteins or using with various numbers of clusters found by the different clustering methods. The results are shown in Table 1. We conclude that the grouping method using city block metrics has the best performance among all conditions for both CGP and CENR in terms of AUC. Although CGP generally has better performance for identifying previously known relationships, it does not perform as well as CENR in terms of accuracy of the model for full stimulation or prediction for B7 data. This may be due to the fact that CGP has few parameters and thus the generalization ability is not as good as CENR.

4 Discussion

There has been extensive modeling of signaling processes in systems such as the T cell (for overview see Chylek *et al.*, 2014), including modeling that considers spatial organization (Angermann *et al.*, 2012). These have been based on known biochemical reactions and thus the need for causal inference is minimal. We have described here a novel approach to understanding the potential causative relationships between the spatiotemporal distributions of proteins involved in signaling or other regulatory networks. We have demonstrated that the approach yields a number of conclusions that are consistent with current knowledge of actin regulators, and also relationships that deserve further exploration.

An important issue is how these methods would scale for larger sets of proteins or more extensive kinetics. For optimizing the CGP, the worst-case time complexity is $O(M^2N^3 + KMN)$ for M blocks in each iteration, where M is the lag time, K is the total number of time points, and N is the number of nodes (Mei and Moura, 2017). In our case, $N = PR$, where P and R are the numbers of proteins and regions. Therefore the time complexity is $O(M^2P^3R^3 + KMPR)$.

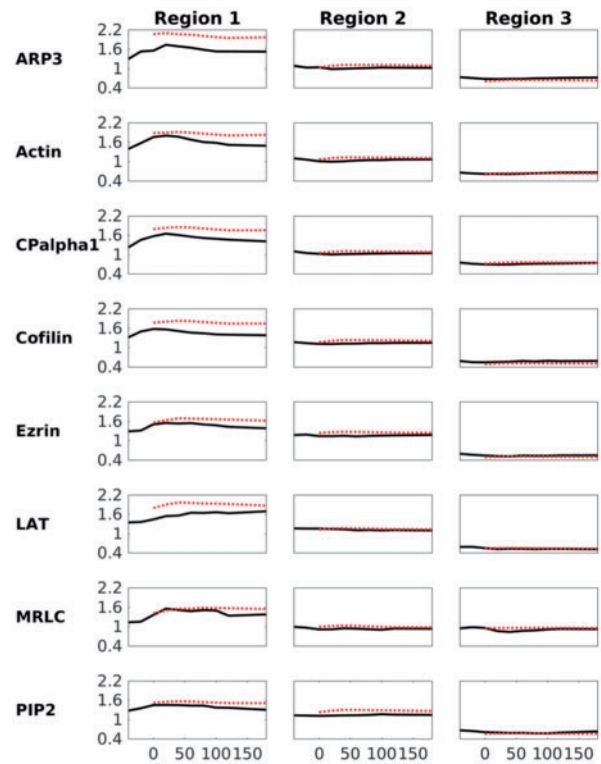


Fig. 8. B7 prediction using CGP. For each figure, the x-axis show time after synapse formation, and the y-axis shows relative concentration of a protein. The actual kinetics are shown in black and the predictions in red

For CENR, the worst-case time complexity is $O(M^2N^3)$ if we can pre-compute some quantities in order to get rid of K for each iteration. Therefore for our problem, the complexity is $O(M^2P^3R^3)$ in each iteration. Though the time complexities are high, we were able to solve them in reasonable time since the numbers of proteins and regions were small. On a single 2.4 GHz Intel node, inference took <10 s for CGP and <1 s for CENR. The full process including cross-validation for parameter estimation and permutation tests for significance took 200 and 25 CPU hours, respectively.

Although the analysis can provide some promising perspective for image-based causal inference for protein networks, we must point out that this analysis is still preliminary and has some limitations. First, the number of proteins is relatively small (12 proteins), due to limited available data. Second, the proteins were not randomly chosen and they therefore may have more relationships than expected. Third, we attempted to remove confounding effects through clustering but this is a complex subject that will require further exploration.

Despite these limitations, we believe that the consistency of the results with prior knowledge, and the fact that the model learned under one condition could make statistically significant predictions of concentrations under another condition not used in the model learning, to be highly encouraging. In the future, we plan to expand the number of proteins analyzed, which will give a clearer picture of the whole network and may identify both additional putative relationships and resolve potentially confounding relationships.

Acknowledgements

We thank Dr Laura McMillan for help in constructing the partial ground truth.

Funding

This work was supported in part by US National Institutes of Health [grant P41 GM103712], by US National Science Foundation [grant MCB1616492] and by UK BBSRC [grant BB/P11578/1-15 NFSBIO].

Conflict of Interest: none declared.

References

- Angermann,B.R. *et al.* (2012) Computational modeling of cellular signaling processes embedded into dynamic spatial contexts. *Nat. Methods*, **9**, 283–289.
- Boyd,S. *et al.* (2011) Distributed optimization and statistical learning via the alternating direction method of multipliers. *Found. Trends Mach. Learn.*, **3**, 1–122.
- Burkhardt,J.K. *et al.* (2008) The actin cytoskeleton in T cell activation. *Annu. Rev. Immunol.*, **26**, 233–259.
- Chang,R. *et al.* (2015) Causal inference in biology networks with integrated belief propagation. In *Pacific Symposium on Biocomputing*, Kohala Coast, Hawaii, p. 359.
- Chindelevitch,L. *et al.* (2012) Causal reasoning on biological networks: interpreting transcriptional changes. *Bioinformatics*, **28**, 1114–1121.
- Chylek,L.A. *et al.* (2014) Modeling biomolecular site dynamics in immunoreceptor signaling systems. In: *A Systems Biology Approach to Blood*. Springer, New York, pp. 245–262.
- Eichler,M. (2012) Causal inference in time series analysis. In: Berzuini,C., Dawid, A.P., and Bernardinelli, L. (eds.) *Causality: Statistical Perspectives and Applications*. John Wiley & Sons, Chichester, pp. 327–354.
- Farkash-Amar,S. *et al.* (2014) Noise genetics: inferring protein function by correlating phenotype with protein levels and localization in individual human cells. *PLoS Genet.*, **10**, e1004176.
- Figueiredo,M.A. *et al.* (2007) Gradient projection for sparse reconstruction: application to compressed sensing and other inverse problems. *IEEE J. Select. Top. Signal Process.*, **1**, 586–597.
- Fujita,A. *et al.* (2010) Granger causality in systems biology: modeling gene networks in time series microarray data using vector autoregressive models. In: *Brazilian Symposium on Bioinformatics*. Springer, Berlin, pp. 13–24.
- Gandhi,M. *et al.* (2009) Coronin switches roles in actin disassembly depending on the nucleotide state of actin. *Mol. Cell*, **34**, 364–374.
- Granger,C.W. (1969) Investigating causal relations by econometric models and cross-spectral methods. *Econometrica*, **37**, 424–438.
- Machacek,M. *et al.* (2009) Coordination of Rho GTPase activities during cell protrusion. *Nature*, **461**, 99–103.
- Mei,J. and Moura,J.M. (2017) Signal processing on graphs: Causal modeling of unstructured data. *IEEE Trans. Signal Process*, **65**, 2077–2092.
- Pearl,J. (2009) *Causality*. Cambridge University Press, Cambridge.
- Roybal,K.T. *et al.* (2016) Computational spatiotemporal analysis identifies WAVE2 and cofilin as joint regulators of costimulation-mediated T cell actin dynamics. *Sci. Signal*, **9**, rs3.
- Segal,E. *et al.* (2003) Module networks: identifying regulatory modules and their condition-specific regulators from gene expression data. *Nat. Genet.*, **34**, 166–176.
- Shojaie,A. and Michailidis,G. (2010) Discovering graphical Granger causality using the truncating lasso penalty. *Bioinformatics*, **26**, i517–i523.
- Urano,T. *et al.* (2003) Haematopoietic lineage cell-specific protein 1 (HS1) promotes actin-related protein (Arp) 2/3 complex-mediated actin polymerization. *Biochem. J.*, **371**, 485–493.
- Vilela,M. *et al.* (2013) Fluctuation analysis of activity biosensor images for the study of information flow in signaling pathways. *Methods Enzymol.*, **519**, 253.
- Welf,E.S. and Danuser,G. (2014) Using fluctuation analysis to establish causal relations between cellular events without experimental perturbation. *Biophys. J.*, **107**, 2492–2498.
- Zou,C. and Feng,J. (2009) Granger causality vs. dynamic Bayesian network inference: a comparative study. *BMC Bioinformatics*, **10**, 122.