PLoS COMPUTATIONAL BIOLOGY

Provisional DOI: 10.1371/journal.pcbi.0010064.eor

Citation: Kaufman A, Keinan A, Meilijson I, Kupiec M, Ruppin E (2005) Quantitative analysis of genetic and neuronal multi-perturbation experiments. PLoS Comput Biol. In press. DOI: 10.1371/journal.pcbi.0010064.eor

Future Article URL: http://dx.doi.org/10.1371/journal.pcbi.0010064

Quantitative Analysis of Genetic and Neuronal Multi-Perturbation Experiments

Alon Kaufman[1,*], Alon Keinan[2,*], Isaac Meilijson[3], Martin Kupiec[4], Eytan Ruppin[2,5]

[1]Center of Neural Computation, Hebrew University, Jerusalem 91904, Israel. [2]School of Computer Science, Tel-Aviv University, Tel-Aviv 69978, Israel. [3]School of Mathematical Sciences, Tel-Aviv University, Tel-Aviv 69978, Israel. [4] Department of Molecular Microbiology and Biotechnology, Tel-Aviv University, Tel-Aviv 69978, Israel. [5]School of Medicine, Tel-Aviv University, Tel-Aviv 69978, Israel. [*]These authors made an equal contribution to this work.

Corresponding author: Eytan Ruppin, School of Computer Science, Tel-Aviv University, Tel-Aviv, 69978, Israel, +972-3-640-6528(voice), +972-3-640-9357(fax), ruppin@post.tau.ac.il

Abstract

Perturbation studies, in which functional performance is measured after deletion, mutation or lesion of elements of a biological system, have been traditionally employed in many fields in biology. The vast majority of these studies have been qualitative and employed single perturbations, often resulting in little phenotypic effect. Recently, new emerging experimental techniques allow to carry out concomitant multi-perturbations and uncover the causal functional contributions of system elements. This study presents the first rigorous and quantitative multi-perturbation analysis of gene knockout and neuronal ablation experiments. In both cases, a quantification of the elements' contributions and new insights and predictions are provided. Multi-perturbation analysis has a potential wide range of applications and is gradually becoming an essential tool in biology.

Synopsis

Which are the important elements of a system? What are their relative contributions to the performance of the various tasks the system is involved in? These simple and basic questions typically arise when coming to analyze the workings of any system in general, and specifically, biological systems. In the latter, the elements may be genes, proteins, cells or tissues, depending on the level and scope of the analysis. To address these questions in a causal manner, perturbations are required, where the elements are perturbed and the resulting performance function is recorded. This approach has been one of the cornerstones of biological research. However, it has been usually confined to the perturbation of a single element at a time, which may lead to misleading results if the elements of the system functionally interact with each other. This paper addresses these questions by providing a quantitative and rigorous method for the analysis of multi-perturbation experiments, where more than one element may be concomitantly perturbed. The workings of the new method are demonstrated in the analysis of genetic multi-knockout experiments of DNA repair in the yeast and a neural circuit in C.elegans accounting for chemotaxis. However, the method is general and can be applied to study many other systems in numerous pertinent biological domains.

Introduction

System identification (localization of function) in biological networks is currently mainly studied in genetics by high-throughput expression profiling and in neuroscience by functional brain imaging. While these techniques have proved to be very useful and productive [1, 2], the correlation-based approach they employ does not necessarily identify causal relations. Previous studies have shown that there may be, at times, a weak correlation between the expression of different genes and their role in various cellular functions [3, 4]. In a similar vein, the need to add a causal perturbation analysis to complement a correlation-based one has also been raised in the neuroscience literature

(*e.g.*, [5, 6]). To causally deduce the roles played by system elements (genes, proteins, neurons, brain regions, etc.), perturbation studies, in which functional performance is measured after deletion, mutation or lesion of the different elements, have been traditionally employed. However, the vast majority of these studies have perturbed only one element at a time, often resulting in little phenotypic effect. Hence, in complex biological systems, multiple concomitant perturbations should be employed to reveal the contributions of the different elements to the system's functioning. In genetics, the lack of phenotypic effects may be due to the existence of duplicates, alternative pathways and functional overlaps [1, 7]. To uncover these effects, new experimental techniques are now emerging to carry out the necessary multi-perturbation studies [8, 9, 1]. Specifically, the recent discovery of RNA interference (RNAi) [10, 11] and the rapid recent advances in gene silencing with RNAi chips [12, 13] may advance multi-perturbation technology to our door step. In neuroscience, techniques as Transcranial Magnetic Stimulation (TMS) allow to induce reversible "virtual lesions", enabling to perform multi-lesion experiments leading to the analysis of cognitive and perceptual tasks in humans [14, 15]. A question remains: how can the results of such multi-perturbation experiments be integrated and analyzed and what knowledge can be extracted from them?

To address this challenge, Keinan et al. [16] have developed the Multi-Perturbation Shapley value Analysis (MPA) method, and presented its application to the analysis of a neural model of the lamprey swimming controller and to the analysis of reversible cooling deactivation experiments in cats. Here we expand these results in two fundamental ways: First, we present the first application of the MPA to the analysis of gene knockout experiments and to the analysis of neuronal ablation data. Second, we present a complementary method for the analysis of multi-perturbation data, the Functional Influence Network (FIN) algorithm. In difference from existing methods for network inference and system identification in biology, which employ methods from machine learning such as Bayesian network inference [17, 18], Boolean networks [19] and other reverse engineering methods [20], both MPA and FIN are based on concepts from Game theory. These new methods are the first to utilize fundamental results from Game Theory to assess the contribution of system elements and functional subsets of such elements to the overall system performance function. For the task of determining the functional contribution of system elements, these game theory tools are more adequate than standard machine learning approaches employing error minimization such as [21], since they are based on a solid axiomatic framework and provide a unique contribution assignment [16]. Recently, there have been a number of attempts to utilize game theory approaches in neuroscience, but these had a completely different goal of constructing decision making models [22, 23].

The goal of MPA is to define and calculate the contribution (importance) of system elements to a certain function, from a data set of a series of multi-perturbation experiments. In each such experiment, a different subset of the system elements is concomitantly perturbed (denoting a perturbation configuration), and the system's performance in the studied function is measured. The FIN algorithm analyzes the same multi-perturbation data. It describes the incremental contribution of each subset of elements to the function studied, and produces a compact representation, composed only of the most important subsets. As the full set of all theoretically possible multi-perturbation

experiments required for the MPA and FIN computation is usually unavailable, both analyses employ a predictor algorithm to compute the system's performance on the missing multi-perturbation experiments.

In the following sections we describe the MPA and FIN methods and present their application to two different biological systems: DNA post replication repair pathway in the yeast *Saccharomyces cerevisiae*, and laser ablation studies of *C. elegans* chemosensory neurons.

Results

Quantitative Multi-Perturbation Analysis

Multi-Perturbation Shapley Value Analysis (MPA)

The starting point of the Multi-Perturbation Shapley value Analysis (MPA) [16] is a data set of multi-perturbation experiments studying a system's performance in a certain function. In each such experiment, a different subset of the system's elements are perturbed concomitantly (denoting a *perturbation configuration*) and the system's performance following the perturbation is measured. Given this data set, the goal of MPA is to ascribe to each element its contribution (importance) in carrying out the studied function.

The basic observation underlying MPA is that the multi-perturbation setup is essentially equivalent to a coalitional game. A *coalitional game* is defined by a pair (N, v), where $N = \{1, \ldots, n\}$ is the set of all *players* and $v(S)$, for every $S \subseteq N$, is a real number associating a *worth* with the *coalition* S, and $v(\emptyset) = 0$. In the context of multi-perturbations, N denotes the set of all the system's elements, and for each $S \subseteq N$, $v(S)$ denotes the performance measured under the perturbation configuration in which all the elements in S are intact and the rest are perturbed.

A *payoff profile* of a coalitional game is the assignment of a payoff to each of the players. A *value* is a function that assigns a unique payoff profile to a coalitional game. It is *efficient* if the sum of the payoffs assigned to all players is $v(N)$. The definite efficient value in game theory and economics for coalitional games is the *Shapley value* [24], defined as follows: let the *marginal importance* of player i to a coalition S, with $i \notin S$, be

$$\Delta_i(S) = v(S \cup \{i\}) - v(S). \tag{1}$$

Then, the Shapley value is defined by the payoff

$$\gamma_i(N, v) = \frac{1}{n!} \sum_{R \in \mathcal{R}} \Delta_i(S_i(R)) \tag{2}$$

assigned to player i, for all $i \in N$, where $\mathcal{R}$ is the set of all $n!$ orderings of N, and $S_i(R)$ is the set of players preceding i in the ordering R. The Shapley value has a clear intuitive interpretation, denoting the average marginal importance of player i to the game. Importantly, it has an axiomatic foundation, which is well suited for the analysis of biological data [16] (see the axiomatic basis of the Shapley Value in Supplementary

Text III). The MPA uses the Shapley value as the unique fair measure of each element's contribution (importance) to the function in question.

Obviously, conducting the large number of multi-perturbation experiments (exponential in the size of the system) required for the computation of the Shapley value is most often intractable. In such cases, MPA involves training a predictor using a given subset of multi-perturbation experiments to predict the performance levels of all missing experiments. Given the predicted outcomes of all multi-perturbation experiments, a *predicted Shapley value* is calculated as the Shapley value based on these outcomes. The accuracy of such an analysis depends on the accuracy of the predictions [16] and is determined using standard cross validation techniques as leave-one-out [25]. The analysis presented throughout this paper, is based on a Projection Pursuit Regression predictor [26]. The accuracy of the predicted contributions is strongly depended on the accuracy of the predictor used. However, the outcome, the predicted contributions, are more accurate then the individual predictions provided by the predictor due to the fact that the Shapley value is obtained via an averaging over a large number of predictions. Assuming that the predictor is unbiased, prediction errors tend to cancel each other out, resulting in a predicted Shapley value which is unbiased and very similar to the real one. See Supplementary Protocol IV for a detailed discussion on the robustness and stability of the predicted contributions produced by the MPA.

Functional Inference Network Analysis (FIN)

The Functional Influence Network (FIN) algorithm, based on the series of multi-perturbation experiments, begins with the computation of a performance prediction function F, in the form of a multi-linear weighted summation over all 2^n subsets (summands) of the n elements in the system (see Methods). The weight of each summand describes the marginal contribution of that subset of elements to the value of F. Given a configuration S of perturbed and intact elements, the goal of $F(S)$ is to accurately describe the experimentally measured performance value of the system in the task studied under that configuration, $v(S)$. After obtaining F, the FIN algorithm proceeds to prune its summands and retain only the most significant ones, to obtain a compact, approximate performance prediction function $\widetilde{F}$ (the detailed algorithm is provided in the Methods section). The latter preserves a high percentage of F's original prediction accuracy but aims to provide a compact functional description that may be visualized, if sufficiently compact. Each summand of F and $\widetilde{F}$ can be viewed as an integrative functional pathway in the sense that the knockout of any of its elements will zero its contribution. In cases where the full set of all possible multi-perturbation experiments required for the FIN computation is unavailable, the FIN, similar to MPA, uses a predictor algorithm to compute the system's performance in the missing experiments. The accuracy of the resulting FIN description for a given prediction accuracy is part of a broader conceptual issue, that of the relationship between "predictive knowledge" (the prediction accuracy) and "descriptive knowledge" (provided in the case in hand by the FIN output). A comprehensive investigation of this fundamental issue is out of the scope of the current work and will be addressed in a separate paper (preliminary results have been recently presented by Kaufman et al. in the BioPathways Special Interest Group, ISMB2005).

Gene Knockout Analysis: The Rad6 Pathway

Post Replication Repair

We performed a multi-knockout study of the DNA Post-Replication Repair (PRR) system of the yeast *Saccharomyces cerevisiae*. DNA repair pathways in yeast have been classified genetically into three major repair systems specialized on different types of damage: (1) The excision repair (Rad3) group, which is mainly involved in the repair of UV-irradiated DNA. (2) The recombination repair (Rad52) group, which is mainly involved in repair of damage caused by ionizing radiation and of double-strand breaks in the DNA, and (3) The post-replication repair (PRR, Rad6) group of genes. This pathway is believed to warrant replication in situations where lesions in the template strand would otherwise cause a stalling of the replication machinery, as occurring following UV radiation.

A key physiological target of the PRR pathway is PCNA, a homotrimeric ring-shaped protein that encircles DNA, functioning as a freely sliding clamp that tethers DNA polymerase to the DNA template. The current hypothesis posits that following the stalling of the replicative DNA polymerases (when lesions are encountered), PCNA is modified, and the replicative polymerase is replaced by trans-lesion polymerases. Ubiquitination of PCNA is carried out by the Rad6 ubiquitin-conjugating enzyme, which is targeted to the stalled replication fork through physical interactions with the Rad18 co-factor [27, 28]. During DNA synthesis, Replication Factor C (RFC), a hetero-pentameric protein complex, is necessary for loading PCNA onto double-stranded DNA at the primer-template junction. Recently, several proteins with similarities to Rfc1 (the large sub unit of RFC) were found to form RFC-like complexes (RLCs), including Elg1(YOR144C/Q12050), Rad24(YER173W/P32641) and Ctf18(YMR078C/P49956). These RLCs may act similarly to RFC, loading PCNA or PCNA-related complexes that act as clamps for specific DNA polymerases [29]. In addition to the three RLC genes, our study includes *RAD18*(YCR066W/P10862), a gene needed for PCNA modification, and the *REV3* (YPL167C/P14284) gene which encodes an alternative DNA polymerase (ζ).

The analyzed data includes a series of multi-knockout experiments carried out in the lab of one of the authors (M.K.), testing the ability of the resulting mutants to resolve the single-stranded gaps created after UV irradiation. Hence, the perturbations are gene knockouts and the elements are the five genes listed above. The performance under investigation is UV survival, measured by the relative number of colonies that survive compared to the wild-type yeast strain (normalized on a scale from 0 to 1). The dataset includes 21 multi-knockout experiments (see Supplementary Table S1). Prediction of the full, 32, multi-knockout set was obtained using Projection Pursuit Regression, explaining 79.6% of the data variance via leave-one-out cross validation.

MPA and FIN Analysis

Figure 1 displays the results of MPA of the Rad6 data, leading to a quantification of the causal contribution of each of the Rad6 genes to PRR repair. The most important genes are *RAD18* and *REV3*, the modifier gene and the DNA polymerase respectively. All three RLCs play a causal role as well, but their importance differs markedly.

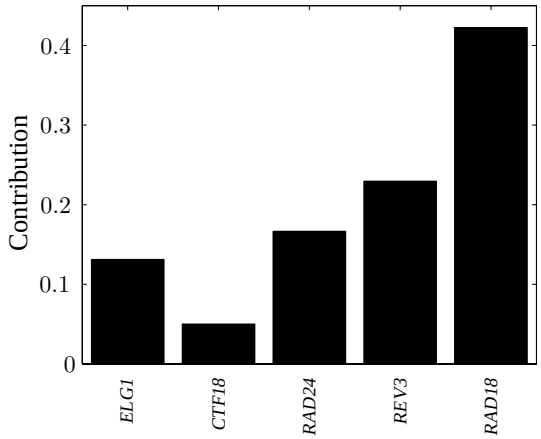

Figure 1: Contributions of genes in the Rad6 pathway to PRR functioning (normalized such that their sum equals 1).

The multi-knockout data can be utilized to construct a unique weighted multi-linear performance prediction function F which, given any configuration of knocked-out and intact genes, can accurately predict the PRR performance level. However, this function contains 32 (or, more generally 2^n) terms corresponding to all possible knockout configurations and hence is unintelligible and uninformative to the biologist. To extract the relevant information in the data and make it explicit, F is further processed via the FIN algorithm to construct a compact and yet fairly accurate functional prediction function, $\widetilde{F}$. Each of the terms in $\widetilde{F}$ can be viewed as a serial functional pathway, whose contribution depends on the intactness of all its component genes[†]. $\widetilde{F}$s compactness can be utilized to visualize the PRR functioning via a FIN diagram, as shown in Figure 2.

The FIN analysis gives rise to two new hypotheses: first, as evident, both *ELG1* and *RAD24* play a significant role even without *REV3* (edges *RAD24–RAD18* and *ELG1–RAD18* in Figure 2). Hence, there is probably another polymerase (or perhaps more than one) involved in the PRR process, suggesting that both *ELG1* and *RAD24* play a role loading this additional DNA polymerase. A good candidate for such an additional polymerase is Pol-η, encoded by the *RAD30* gene. This alternative polymerase has been shown to be dependent on Rad6/Rad18 for activity. Second, the *REV3–RAD18* pathway (edge *REV3–RAD18* in Figure 2) encompasses 26% of the system's repair performance. This suggests that there are some additional DNA polymerase loaders beside those investigated. Alternatively, some of the functionality of the DNA polymerase may be maintained even in the absence of the RLCs.

[†]Obviously, membership in a functional pathway does not necessarily imply that there are direct physical interactions between the elements of the pathway.

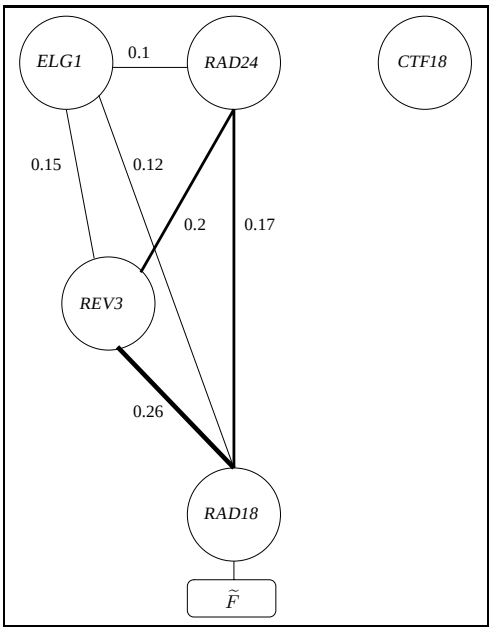

Figure 2: The FIN diagram of the PRR pathway visualizes the compact performance function $\widetilde{F} = e \cdot (d \cdot 0.26 + c \cdot 0.17 + a \cdot 0.12 + d \cdot (c \cdot 0.2 + a \cdot 0.15) + (c \cdot a) \cdot 0.1)$, where $a...e$ are Boolean variables representing the genes(a =*ELG1* , b =*CTF18* , c =*RAD24* , d =*REV3* and e =*RAD18*). The investigated genes are represented as binary nodes, whose state is determined according to the state of the corresponding genes in a given perturbation configuration, intact or knocked-out. The nodes are connected with edges, their weight representing the functional influence between the two endpoint genes (the width of the edge is proportional to its weight). Given a knockout configuration, the expected performance level $\widetilde{F}$ can be calculated by summing up the weights on the edges between intact nodes which form a connected component with the function node $\widetilde{F}$. For example, in a mutant where both *REV3* and *ELG1* are knocked-out, the intact nodes are *CTF18, RAD24* and *RAD18.* The edge *RAD24-RAD18* is the resulting connected subgraph of $\widetilde{F}$, predicting a performance level of 0.17. Observe that there are three main (two-node) pathways leading to $\widetilde{F}$, (edges *RAD24–RAD18* , *ELG1–RAD18* and *REV3–RAD18*), where *RAD18* is an essential gene in all of them. The RLC *CTF18* has no significance in the FIN description even though it has a contribution of 4% (Figure 1); it contributes marginally across many insignificant summands and does not play a significant role in any major one.

Analysis of Neuronal Ablations: Chemotaxis in C. elegans

We turn to address the question of function localization in the nervous system, focusing on laser ablation experiments of the C. elegans chemosensory neurons. The behavior studied in these experiments is Chemotaxis, in which the nematode directs its movements according to chemical gradients in the environment, moving toward the highest concentration of food or fleeing from toxins. Re-analyzing the data published by Bargmann and Horvitz [30], we compare the qualitative conclusions given in their paper to the quantitative analogues obtained by applying the MPA. The elements studied are 8 sensory neuron pairs (out of a total of 16 pairs which form the chemosensory system [31]). In each laser ablation experiment both neurons in a pair are either intact or perturbed. The performance measures, chemotaxis to various attractants (each composing a distinct functional task), were evaluated by placing the animal on an agar plate with a gradient of an attractant on one side of the plate, and scoring the chemotaxis performance by counting the number of times the animal arrived at the peak of the gradient minus the number of times the animal arrived at the control plug at the opposite side of the plate. The level of chemotaxis performance was evaluated under 31 perturbation configurations, according to the protocol described in [30]. Prediction of the full, 256 multi-lesion set needed to calculate the neurons' contributions was obtained using Projection Pursuit Regression as the predictor. A cross validation leave-one-out procedure shows that the predictor explains $65-80\%$ of the data variance depending on the attractant type.

Neuronal Contribution Analysis

Figure 3 displays the contributions of the different neuron pairs to 4 different attractants tasks (Serotonin, Cl, cAMP and Biotin). As evident, the ASE pair is the most important to chemotaxis across all attractants, in line with the results of Bargmann and Horvitz. However, MPA additionally shows how the importance of all other neurons varies among the different attractants. The processing of the Serotonin task is more distributed (the neuronal contributions are more equally spread) across the network than that of the other tasks. A FIN diagram of the relative simple and localized cAMP task is provided in Supplementary Protocol V. Notably, the ASH pair has a negative, inhibitory, contribution to chemotaxis (chemotaxis will be more successful on average if the ASH pair is ablated). This observation, is in line with other more recent experimental assays showing that ASH plays a role in mediating C. elegans avoidance of toxic chemicals [32]. Interestingly, examining the interaction between these two neuron pairs, ASE and ASH show a negative interaction: In three attractant tasks, Serotonin, Biotin and cAMP. ASE's contribution is suppressed by ASH, i.e., the contribution of ASE is lower when ASH is intact than when it is ablated, and similarly the ASH is suppressed by ASE. Thus, predicting that ASE will be an antagonist of avoidance behavior where ASH is likely to be highly activated.

Multiple Tasks Analysis

The contributions of the neuron pairs across the different tasks can be summarized in a contribution matrix, where C_{ij} in the matrix denotes the contribution of element j

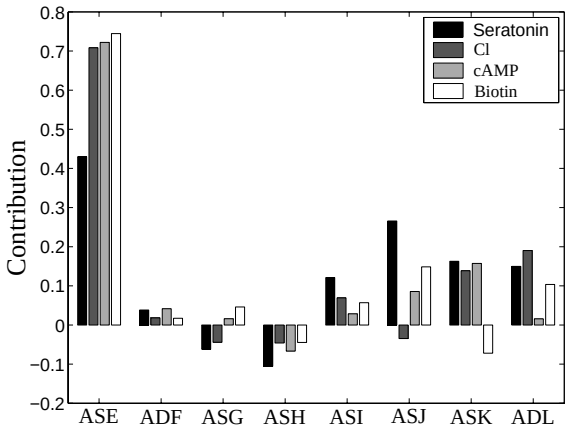

Figure 3: Contributions of the 8 neuron pairs to the different chemotaxis attractant tasks (normalized such that their sum for each attractant equals 1).

to task i (Supplementary Table S2). This matrix description permits the utilization of a series of analyses which are not applicable when the results are only summarized in a qualitative manner. Singular Value Decomposition (SVD), a standard method for dimension reduction previously utilized in various biological applications (e.g., [33]), can be applied to the contribution matrix to reveal both its "neuron-space" and its "task-space", identifying similar tasks and similar functional contributions of neurons. Figure 4A presents the results of an SVD of the contribution matrix in the task space. The figure shows four main clusters of the neurons, based on the contributions of different neurons across the attractants (i.e., the column vectors of the contribution matrix in Supplementary Table S2). The distinct placement of the ASE and ASJ neurons is notable. Neurons participating in each of the clusters $\{$ADF,ASG,ASH$\}$ and $\{$ASI,ASK,ADL$\}$ have similar functional roles across the investigated attractants. Figure 4B presents the results of SVD in the neuron space. The processing of Serotonin chemotaxis is localized very differently than the processing of the other attractants, and the processing of cAMP and Biotin is localized in a very similar manner. The similarity between the tasks was already observed in [30], although not shown in a rigorous manner. Chemotaxis to Cl, which was thought to be processed similarly to that of cAMP and Biotin [30], is actually processed quite differently.

Discussion

This paper presents a multi-perturbation analysis of two different biological systems. The analysis reinforces previous known knowledge in a quantitative manner and leads to new insights. The MPA analysis of the PRR system shows that each of the RLCs has a different magnitude of contribution to the PRR process. The FIN analysis gives rise

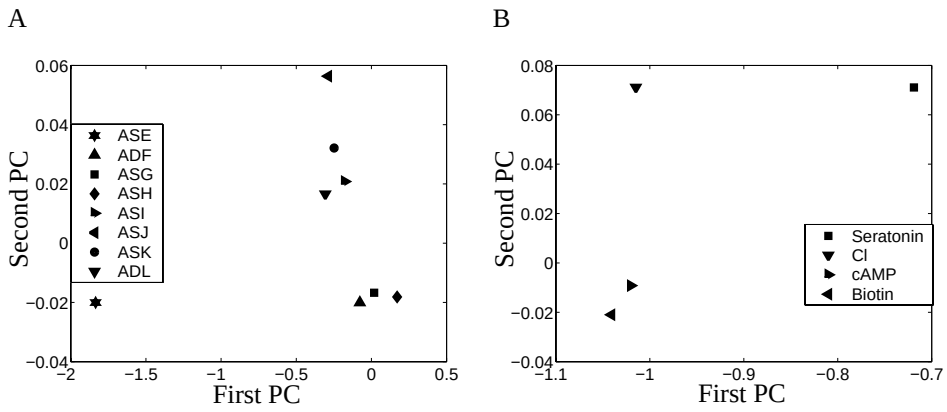

Figure 4: SVD analysis of the contribution matrix, using the two main principal components of the SVD decomposition, which together explain 96% of the data's variance. A). *"Task-space"*, presenting the projections of the neurons contribution vectors (columns vectors of the contribution matrix) onto the two main principal eigenvectors of the task space. B). *"Neuron-space"*, presenting the projections of the tasks contribution vectors (row vectors of the contribution matrix) onto the two main principal eigenvectors of the neuron space.

to the hypotheses that there are additional polymerase loading complexes in the yeast, and that DNA polymerase ζ encoded by *REV3* is probably not the only polymerase involved in PRR. The analysis of C. elegans' chemotaxis provides a more refined picture of the sensory network and rigorously reinforces previous findings.

MPA and FIN are the first methods to harness game theory concepts for the analysis of biological systems. Further work is needed to better adapt these methods to the constraints of biological systems, most notably, the limited depth (i.e., number of concomitantly perturbed elements) of multi-perturbations in biology. However, this is likely to be a very rewarding endeavor, as such multi-perturbation analysis has potentially many applications. The most direct and natural ones are those concerning the analysis of causal perturbation data; in genetics, using gene silencing with RNAi. In neuroscience, there is now a new prospect of carrying out experimental perturbation studies using Transcranial Magnetic Stimulation (TMS). This technique allows to induce "virtual lesions" in normal subjects performing various cognitive and perceptual tasks [14, 15].

Importantly, MPA and FIN are not limited to causal perturbation analysis, where one controls the lesions made. They may well be applied to sets of naturally given multi-perturbations, e.g., by studying the brain localization of cognitive functions from pertaining "multi-lesion" data in stroke patients. In summary, multi-perturbation studies are a necessity if one wants to understand the processing of biological networks in a quantitatively causal manner. The methods described in this paper are a harbinger of this new kind of studies, offering a novel and rigorous way of making sense out of

them.

Methods

The basic MPA and FIN analysis methods are described at the beginning of the Results section. Here we provide a description of the extension of MPA to a 2-dimensional interaction analysis and the details of the FIN algorithm.

MPA Interaction Analysis

In complex systems, the importance of an element may strongly depend on the state (perturbed or intact) of other elements. A higher order description may be necessary to capture these interactions. Such *high-dimensional analysis* provides further insights into the network's functional organization.

We focus on the description of two-dimensional interactions. A natural definition of the latter is as follows [16]: let $\gamma_{i,\bar{j}} = \gamma_i(N \setminus \{j\}, v^{N\setminus\{j\}})$ be the Shapley value of element i in the subgame of all elements without element j, where $v^{N\setminus\{j\}}$ is the value function over the set $(N \setminus \{j\})$ which denotes the set N without the element j. Intuitively, this is the average marginal importance of element i when element j is perturbed.

Let us now define the coalitional game (M, v^M), where $M = N \setminus \{i,j\} \cup \{(i,j)\}$ ((i,j) is a new compound element composed of both i and j) and $v^M(S)$, for $S \subseteq M$, is defined by

$$v^M(S) = \begin{cases} v(S) & : \quad (i,j) \notin S \\ v(S \setminus \{(i,j)\} \cup \{i,j\}) & : \quad (i,j) \in S \end{cases} \tag{3}$$

where v is the payoff function of the original game with elements N. Then $\gamma_{(i,j)} = \gamma_{(i,j)}(M, v^M)$, the Shapley value of element (i,j), is the average marginal importance of elements i and j when jointly added to a configuration. The two-dimensional interaction between element i and element j, $j \neq i$, is then defined as

$$I_{i,j} = \gamma_{(i,j)} - \gamma_{i,\bar{j}} - \gamma_{j,\bar{i}}\,, \tag{4}$$

which quantifies how much the average marginal importance of the two elements together is larger (or smaller) than the sum of the average marginal importance of each of them when the other is perturbed. Intuitively, this symmetric definition ($I_{i,j} = I_{j,i}$) quantifies the synergistic interaction between elements i and j, denoting how much "the whole is greater than the sum of its parts". In cases where the whole is smaller than the sum of its parts, that is, when the two elements exhibit functional overlap or redundancy, the interaction is negative. Based on the genetic interaction nomenclature of Brendel et al. [34], an interaction will be defined as *Epistatic* if $\gamma_{i,\bar{j}}$ is zero and $\gamma_{i,j}$ has a positive contribution, i.e., the intactness of j is essential for i's contribution. The Shapley interaction index [35] provides a more general measure for the interaction among players.

A Detailed Description of the FIN Analysis

Computing F: The Performance Prediction Function

The performance prediction function $F(S)$ can be uniquely computed as the sum $\sum_{T \subseteq S} a(T)$, where S denotes the subset of intact elements in a given perturbation configuration [36, 37]. The coefficients $a(T)$ of the summands are the *dividends*, describing the incremental importance of each summand T to the performance being studied. These dividends can be uniquely calculated from the multi-perturbation data (both given and predicted) according to Eq.(5) (the cardinality of the sets S and T are denoted by corresponding lower cases $s = |S|$,$t = |T|$),

$$a(S) = \sum_{T \subseteq S} (-1)^{t-s} v(T), \quad \forall S \subseteq N. \tag{5}$$

The dividend computation is performed in an iterated manner. It begins from the dividend of the null group and each iteration computes the dividend (incremental contribution) of subsequently larger, subsuming subsets.

Computing $\widetilde{F}$: The Compact Performance Prediction Function

To compute a compact and intelligible approximation of F, a greedy heuristic algorithm is employed which retains only the summands T with the largest dividends $a(T)$, while maintaining a pre-defined level of prediction accuracy. The latter is measured with respect to the performance of the original F (by the normalized Mean Squared Error between $\widetilde{F}$ and F over all perturbation configurations). The algorithm first selects statistically significant summands (based on a null hypothesis that the dividend magnitude is zero), and then eliminates those with a low magnitude to obtain $\widetilde{F}$. To visualize $\widetilde{F}$, we construct the FIN diagram. This construction starts with an algebraic simplification, rewriting $\widetilde{F}$ to minimize the number of appearances of each element. This is done by combining clauses and placing elements common to a few summands as multipliers of the weighted summation of the corresponding, residual summands. In the DNA post replication repair investigation (in the Results section), for example, this stage results in the function $\widetilde{F} = e \cdot (d \cdot 0.26 + c \cdot 0.17 + a \cdot 0.12 + d \cdot (c \cdot 0.2 + a \cdot 0.15) + (c \cdot a) \cdot 0.1)$, where $a...e$ are Boolean variables representing the genes, assigned 1 if the gene is intact and 0 if it is knocked-out. Based on this simplified representation, we construct the FIN diagram by starting from the function node, $\widetilde{F}$, and connecting it to the variables at the most external level parentheses, assigning weights to the connections according to the corresponding dividend coefficients. This process is recursively repeated by connecting the current leaf nodes on each pathway from the node $\widetilde{F}$ to the next level of elements in the remaining parentheses, until the nodes at the most internal parentheses are connected. The resulting FIN diagram depicts the most important functional pathways (interconnected subsets of elements whose contribution depends on the intactness of the other elements in the pathway) and quantifies their relative importance to the function in hand (see Figure 2).

Accession Numbers

Accession numbers for S. cerevisiae proteins are from the SwissProt database (http://www.ebi.ac.uk/swissprot/): Rad18(P10862), Rev3(P14284), Rad24(P32641), Elg1(Q12050), Ctf18(P49956). The accession numbers for genes are from SGD (http://genome-www.stanford.edu/Saccharomyces/) (ORF/SGD identification number): RAD18 (YCR066W/S000662), REV3(YPL167C/S006088), RAD24(YER173W/S000975), ELG1 (YOR144C/S005670), CTF18(YMR078C/S004683).

Acknowledgment

We thank Nir Yosef, Ya'acov Ritov, Shawn Lockery and Cori Bargmann for their valuable comments and suggestions. A. Kaufman is supported by the Yeshaya Horowitz Association through the Center of Complexity Science. A. Keinan was supported by the Dan David Prize Scholarship. M. Kupiec was supported by grants from the Recanati Foundation and the Israel Cancer Fund. E. Ruppin's research is supported by the Tauber Fund and the Center of Complexity Science.

References

[1] Steinmetz L, Davis R (2004) Maximizing the potential of functional genomics. Nat Rev Genet 5:190–201.

[2] Friston K (2002) Beyond phrenology: What can neuroimaging tell us about distributed circuitry? Annual Review of Neuroscience 25:221–250–47.

[3] Winzeler E, Shoemaker D, Astromoff A, Liang H, Anderson K, et al. (1999) Functional characterization of the S. cerevisiae genome by gene deletion and parallel analysis. Science 285:901–906.

[4] Giaever G, Chu A, Ni L, Connelly C, Riles L, et al. (2002) Functional profiling of the Saccharomyces cerevisiae genome. Nature 418 418:387–391.

[5] Kosslyn S (1999) If neuroimaging is the answer, what is the question? Philosophical Transactions of the Royal Society of London: Series B 354:1283–1294.

[6] Grobstein P (1990) Strategies for analyzing complex organization in the nervous system: I. lesion experiments. In: Schwartz E, editor, Computational neuroscience, Cambridge, MA: MIT Press.

[7] Gu Z, Steinmetz L, Gu X, Scharfe C, Davis R, et al. (2003) Role of duplicate genes in genetic robustness against null mutations. Nature 421:63–6.

[8] Carpenter A, Sabatini D (2004) Systematic genome-wide screens of gene function. Nat Rev Genet 5:11–22.

[9] Tong A, Lesage G, Bader G, Ding H, Xu H, et al. (2004) Global mapping of the yeast genetic interaction network. Science 303:808–813.

[10] Hammond S, Caudy A, Hannon G (2001) Post-transcriptional gene silencing by double-stranded RNA. Nature Rev Gen 2:110–9.

[11] Couzin J (2002) Small RNAs make big splash. Science 298:2296–7.

[12] Wheeler D, Bailey S, Guertin D, Carpenter A, Higgins C, et al. (2004) RNAi living-cell microarrays for loss-of-function screens in drosophila melanogaster cells. Nature Methods 1:127–132.

[13] Lehner B, Fraser A (2004) 5,000 RNAi experiments on a chip. Nature Methods 1:103–104.

[14] Pascual-Leone A, Wasserman E, Davey N, Rothwell J (2002) Handbook of Transcranial Magnetic Stimulation. Oxford University Press.

[15] Rafal R (2001) Virtual neurology. Nature Neuroscience 4:862–864.

[16] Keinan A, Sandbank B, Hilgetag C, Meilijson I, Ruppin E (2004) Fair attribution of functional contribution in artificial and biological networks. Neural Comput 16:1887–1915.

[17] Friedman N (2004) Inferring cellular networks using probabilistic graphical model. Science 303:799–805.

[18] Pe'er D, Regev A, Elidan G, Friedman N (2001) Inferring subnetworks from perturbed expression profiles. Bioinformatics 17.

[19] Ideker T, Thorsson V, Karp R (2000) Discovery of regulatory interactions through perturbation: inference and experimental design. In Proc of the Pacic Symp on Biocomputing :305–16.

[20] Tegner J, Yeung M, Hasty J, Collins J (2003) Reverse engineering gene networks: Integrating genetic perturbations with dynamical modeling. Proc Natl Acad Sci USA 100:5944–49.

[21] Aharonov R, Segev L, Meilijson I, Ruppin E (2003) Localization of function via lesion analysis. Neural Computation 15:885–913.

[22] Glimcher P (2002) Decisions, decisions, decisions: choosing a biological science of choice. Neuron 36:323–32.

[23] Braeutigam S (2005) Neuroeconomicsfrom neural systems to economic behaviour. Brain Research Bulletin, In Press .

[24] Shapley L (1953) A value for n-person games. In: Kuhn HW, Tucker AW, editors, Contributions to the Theory of Games, Princeton: Princeton University Press, volume II of *Annals of Mathematics Studies 28*. pp. 307–317.

[25] Stone M (1974) Cross-validatory choice and assessment of statistical predictions. J Royal Stat Soc B 36:111–47.

[26] Friedman J, Stuetzle W (1981) Projection pursuit regression. J Amer Statist Assoc 76:817–823.

[27] Barbour L, Xiao W (2003) Regulation of alternative replication bypass pathways at stalled replication forks and its effects on genome stability: a yeast model. Mutat Res 532:137–155.

[28] Johnson R, Prakash S, Prakash L (1999) Efficient bypass of a thymine-thymine dimer by yeast DNA polymerase, poleta. Science 283:1001–4.

[29] Ben-Aroya S, Koren A, Liefshitz B, Steinlauf R, Kupiec M (2003) *ELG1*, a yeast gene required for genome stability, forms a complex related to replication factor C. Proc Natl Acad Sci USA 100:9906–11.

[30] Bargmann C, Horvitz H (1991) Chemosensory neurons with overlapping functions direct chemotaxis to multiple chemicals in C.elegans. Neuron 7:729–742.

[31] Bargmann C, Mori I (1997) Chemotaxis and thermotaxis. In C. elegans, Cold Spring Harbor Press. II edition, pp. 717–737.

[32] Hilliard M, Bargmann C, Bazzicalupo P (2002) C.elegans responds to chemical repellents by integrating sensory inputs from the head and the tail. Curr Biol 12:730–4.

[33] Alter O, Brown P, Botstein D (2000) Singular value decomposition for genome-wide expression data processing and modeling. Proc Natl Acad Sci USA 97:10101–6.

[34] Brendel M, Haynes R (1973) Interactions among genes controlling sensitivity to radiation and alkylation in yeast. Mol Gen Genet 125:197–216.

[35] Grabisch M, Roubens M (1999) An axiomatic approach to the concept of interaction among players in cooperative games. International Journal of Game Theory 28:547–565.

[36] Grabisch M, Marichal J, Roubens M (2000) Equivalent representations of a set function with applications to game theory and multicriteria decision making. Mathematics of Operations Research 25:157–178.

[37] Harsanyi J (1963) A simplified bargaining model for the n-person cooperative game. International Economic Review 4:194–220.

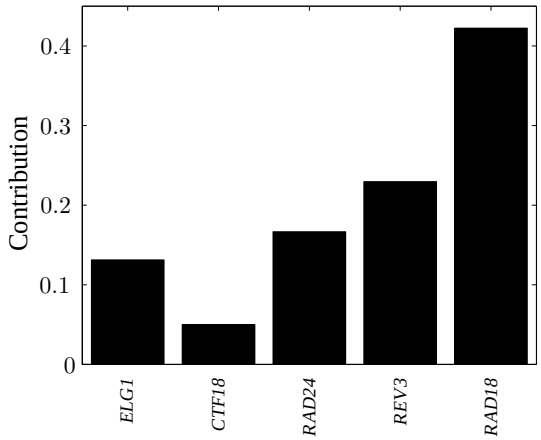

Figure 1: Contributions of genes in the Rad6 pathway to PRR functioning (normalized such that their sum equals 1).

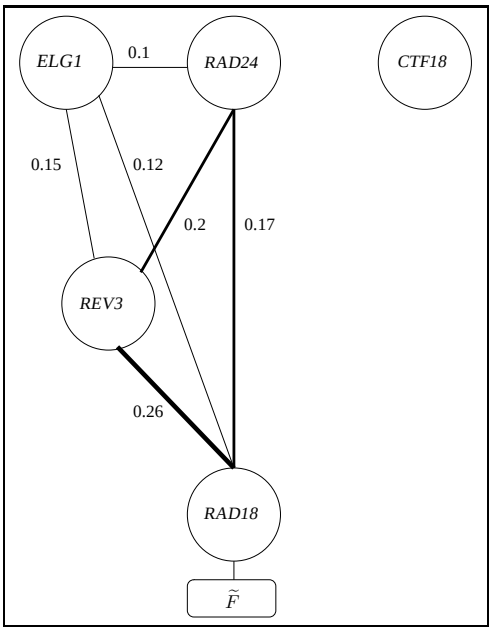

Figure 2: The FIN diagram of the PRR pathway visualizes the compact performance function $\widetilde{F} = e \cdot (d \cdot 0.26 + c \cdot 0.17 + a \cdot 0.12 + d \cdot (c \cdot 0.2 + a \cdot 0.15) + (c \cdot a) \cdot 0.1)$, where $a...e$ are Boolean variables representing the genes(a =*ELG1* , b =*CTF18* , c =*RAD24* , d =*REV3* and e =*RAD18*). The investigated genes are represented as binary nodes, whose state is determined according to the state of the corresponding genes in a given perturbation configuration, intact or knocked-out. The nodes are connected with edges, their weight representing the functional influence between the two endpoint genes (the width of the edge is proportional to its weight). Given a knockout configuration, the expected performance level $\widetilde{F}$ can be calculated by summing up the weights on the edges between intact nodes which form a connected component with the function node $\widetilde{F}$. For example, in a mutant where both *REV3* and *ELG1* are knocked-out, the intact nodes are *CTF18, RAD24* and *RAD18.* The edge *RAD24-RAD18* is the resulting connected subgraph of $\widetilde{F}$, predicting a performance level of 0.17. Observe that there are three main (two-node) pathways leading to $\widetilde{F}$, (edges *RAD24–RAD18* , *ELG1–RAD18* and *REV3–RAD18*), where *RAD18* is an essential gene in all of them. The RLC *CTF18* has no significance in the FIN description even though it has a contribution of 4% (Figure 1); it contributes marginally across many insignificant summands and does not play a significant role in any major one.

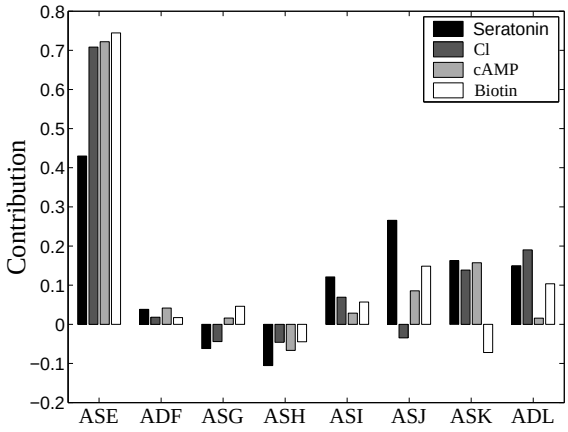

Figure 3: Contributions of the 8 neuron pairs to the different chemotaxis attractant tasks (normalized such that their sum for each attractant equals 1).

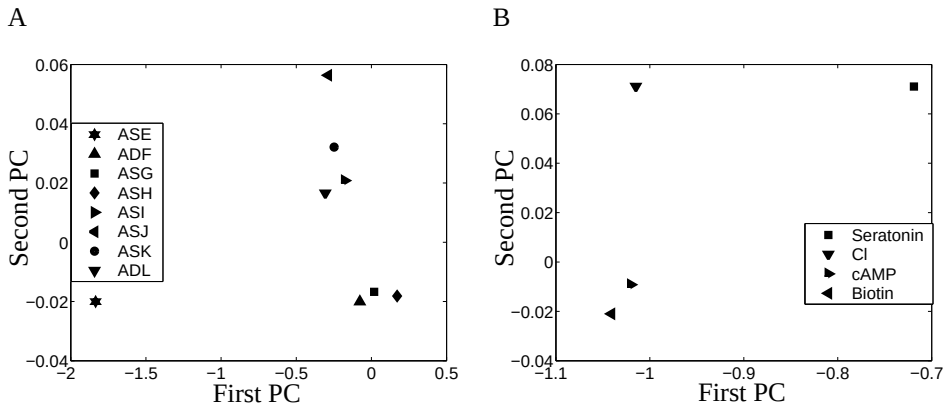

Figure 4: SVD analysis of the contribution matrix, using the two main principal components of the SVD decomposition, which together explain 96% of the data's variance. A). *"Task-space"*, presenting the projections of the neurons contribution vectors (columns vectors of the contribution matrix) onto the two main principal eigenvectors of the task space. B). *"Neuron-space"*, presenting the projections of the tasks contribution vectors (row vectors of the contribution matrix) onto the two main principal eigenvectors of the neuron space.