

MIT Open Access Articles

Wisdom of crowds for robust gene network inference

The MIT Faculty has made this article openly available. **Please share** how this access benefits you. Your story matters.

Citation: Marbach, Daniel, James C Costello, Robert Küffner, Nicole M Vega, Robert J Prill, Diogo M Camacho, Kyle R Allison, et al. "Wisdom of Crowds for Robust Gene Network Inference." *Nature Methods* 9, no. 8 (July 15, 2012): 796–804.

As Published: <http://dx.doi.org/10.1038/nmeth.2016>

Publisher: Nature Publishing Group

Persistent URL: <http://hdl.handle.net/1721.1/87028>

Version: Author's final manuscript: final author's manuscript post peer review, without publisher's formatting or copy editing

Terms of Use: Article is made available in accordance with the publisher's policy and may be subject to US copyright law. Please refer to the publisher's site for terms of use.





Published in final edited form as:

Nat Methods. ; 9(8): 796–804. doi:10.1038/nmeth.2016.

Wisdom of crowds for robust gene network inference

Daniel Marbach^{1,2,8}, James C. Costello^{3,8}, Robert Küffner^{4,8}, Nicci Vega³, Robert J. Prill⁵, Diogo M. Camacho³, Kyle R. Allison³, the DREAM5 Consortium⁶, Manolis Kellis^{1,2}, James J. Collins³, and Gustavo Stolovitzky^{5,7}

¹Computer Science and Artificial Intelligence Laboratory (CSAIL), Massachusetts Institute of Technology (MIT), Cambridge, Massachusetts, USA

²Broad Institute of MIT and Harvard, Cambridge, Massachusetts, USA

³Howard Hughes Medical Institute, Department of Biomedical Engineering, and Center for BioDynamics, Boston University, Boston, Massachusetts, USA

⁴Ludwig-Maximilians University, Department of Informatics, Munich, Germany

⁵IBM T. J. Watson Research Center, Yorktown Heights, New York, USA

Abstract

Reconstructing gene regulatory networks from high-throughput data is a long-standing problem. Through the DREAM project (Dialogue on Reverse Engineering Assessment and Methods), we performed a comprehensive blind assessment of over thirty network inference methods on *Escherichia coli*, *Staphylococcus aureus*, *Saccharomyces cerevisiae*, and *in silico* microarray data. We characterize performance, data requirements, and inherent biases of different inference approaches offering guidelines for both algorithm application and development. We observe that no single inference method performs optimally across all datasets. In contrast, integration of predictions from multiple inference methods shows robust and high performance across diverse datasets. Thereby, we construct high-confidence networks for *E. coli* and *S. aureus*, each comprising ~1700 transcriptional interactions at an estimated precision of 50%. We experimentally test 53 novel interactions in *E. coli*, of which 23 were supported (43%). Our results establish community-based methods as a powerful and robust tool for the inference of transcriptional gene regulatory networks.

Introduction

“The wisdom of crowds,” refers to the phenomenon in which the collective knowledge of a community is greater than the knowledge of any individual¹. Based on this concept, we developed a community approach to address one of the long-standing challenges in molecular and computational biology, which is to uncover and model gene regulatory networks. Genome-scale inference of transcriptional gene regulation has become possible with the advent of high-throughput technologies such as microarrays and RNA sequencing, as they provide snapshots of the transcriptome under many tested experimental conditions.

⁷Correspondence should be addressed to G.S. (gustavo@us.ibm.com).

⁶The complete list of contributors appears at the end of the paper.

⁸These authors contributed equally to this work.

Author contributions

D.M., J.C.C., D.M.C., R.J.P., M.K., J.J.C., and G.S. conceived the challenge; R.J.P. and G.S. performed team scoring; N.V. and K.R.A. performed experimental validation; D.M., J.C.C., R.K., R.J.P., and G.S. performed research; D.M., J.C.C., R.K., N.V., R.J.P., K.R.A., M.K., J.J.C., and G.S. analyzed results; D.M., J.C.C., R.K., M.K., J.J.C., and G.S. wrote the paper; and challenge participants performed network inference and provided method descriptions.

From these data, the challenge is to computationally predict direct regulatory interactions between a transcription factor and its target genes; the aggregate of all predicted interactions comprise the gene regulatory network. A wide range of network inference methods have been developed to address this challenge, from those exclusive to gene expression data^{2,3} to methods that integrate multiple classes of data⁴⁻⁷. These approaches have been successfully used to address many biological problems⁸⁻¹¹, yet when applied to the same data, they can generate quite disparate sets of predicted interactions^{2,3}.

Understanding the advantages and limitations of different network inference methods is critical for their effective application in a given biological context. The DREAM project has been established as a framework to enable such an assessment through standardized performance metrics and common benchmarks¹² (www.the-dream-project.org). DREAM is organized around annual challenges, whereby the community of network inference experts is solicited to run their algorithms on benchmark datasets, participating teams submit their solutions to the challenge, and the submissions are evaluated¹²⁻¹⁴.

Here, we present the results for the transcriptional network inference challenge from DREAM5, the fifth annual set of DREAM systems biology challenges. The community of network inference experts was invited to infer genome-scale transcriptional regulatory networks from gene expression microarray datasets for a prokaryotic model organism (*E. coli*), a eukaryotic model organism (*S. cerevisiae*), a human pathogen (*S. aureus*), as well as an *in silico* benchmark (Fig. 1).

The predictions made from this challenge enable the first comprehensive characterization of network inference methods across different species and datasets, providing insights into method performance, data requirements, and inherent biases. We find that the performance of inference methods varies strongly, with a different method performing best in each setting. Taking advantage of variation, we integrate predictions across inference methods and demonstrate that the resulting community-based consensus networks are robust across species and datasets, achieving by far the best overall performance. Finally, we construct high-confidence consensus networks for *E. coli* and *S. aureus*, and experimentally test novel regulatory interactions in *E. coli*.

We make all benchmark datasets and team predictions, along with the integrated community predictions available as a public resource (Supplementary Data 1–5). In addition, we provide a web interface through the GenePattern genomic analysis platform¹⁵ (GP-DREAM, <http://dream.broadinstitute.org>), which allows researchers to apply top performing inference methods and construct consensus networks.

Results

Network inference methods

Based on the DREAM5 challenge (Supplementary Notes 1–3), we compared 35 individual methods for inference of gene regulatory networks: 29 submitted by participants and an additional 6 commonly used “off-the-shelf” tools (Table 1). Based on descriptions provided by participants, the methods were classified into six categories: Regression, Mutual information, Correlation, Bayesian networks, Meta (methods that combine several different approaches), and Other (methods that do not belong to any of the previous categories) (Table 1).

Performance of network inference methods

We used three gold standards for performance evaluation: experimentally validated interactions from a curated database (RegulonDB¹⁶) for *E. coli*; a high-confidence set of

interactions supported by genome-wide transcription factor binding data¹⁷ (ChIP-chip) and evolutionarily conserved binding motifs¹⁸ for *S. cerevisiae*; and the known network for the *in silico* dataset (**Methods**). Performance on *S. aureus* was evaluated separately (see below) as there currently does not exist a sufficiently large set of experimentally validated interactions.

We assessed method performance for the *E. coli*, *S. cerevisiae*, and *in silico* datasets using the area under the precision-recall (AUPR) and receiver operating characteristic (AUROC) curves¹⁴, and an overall score that summarizes the performance across the three networks (**Methods** and Supplementary Note 4). Figure 2a shows the overall score and the performance on each network for all applied inference methods. On average, regulatory interactions were recovered much more reliably for the *in silico* and *E. coli* datasets compared to *S. cerevisiae*.

Interestingly, well-established “off-the-shelf” inference methods, such as CLR¹¹ and ARACNE⁹ (Mutual Information 1 and 3), were significantly outperformed by several teams. The two teams with the best overall score used novel inference approaches based on random forests¹⁹ and ANOVA²⁰ (Other 1 and 2), respectively (Table 1). However, when considering the performance on individual networks, these two inference methods only performed best for *E. coli*. Two regression methods achieved the best AUPR for the *in silico* benchmark (Regression 1 and 2) and two meta predictors for *S. cerevisiae* (Meta 1 and 5).

There was also strong variation of performance within each category of inference methods (Fig. 2a). For example, the overall scores obtained by regression methods range from the third best of the challenge, down to the fourth lowest. A similar spread in performance can be observed for other categories. We conclude that there is no superior category of inference methods and that performance depends largely on the specific implementation of each individual method. For example, several inference methods used the same sparse linear regression approach (lasso²¹), but exhibited large variation in performance because they implemented different data resampling strategies (Table 1 and Fig. 2a).

Complementarity of different inference methods

To examine the observed variation in performance, we analyzed complementary advantages and limitations of the different methods. As a first step, we explored the predicted interactions of all assessed methods by principal component analysis (**Methods**). The top principal components reveal four clusters of inference methods, which coincide with the major categories of inference approaches (Fig. 2b). Even though the prediction accuracy of methods from the same category varied strongly (Fig. 2a), PCA revealed they have an intrinsic bias to predict similar interactions.

We next analyzed how method-specific biases influenced the recovery of different connectivity patterns (network motifs), which revealed characteristic trends for different method categories (Fig. 2c). For example, feed-forward loops were most reliably recovered by mutual information and correlation-based methods, whereas sparse regression and Bayesian network methods performed worse at this task. The reason for this is the latter approaches preferentially select regulators that independently contribute to the expression of target genes. However, the assumption of independence is violated for genes regulated by mutually dependent transcription factors, as in the case of feed-forward loops. Indeed, linear cascades were more accurately predicted by regression and Bayesian network methods. This shows that current methods trade performance on cascades for performance on feed-forward loops (or vice versa).

For a subset of the transcription factors contained in the gold standards, knockout or overexpression experiments were supplied to DREAM5 participants, and a number of inference methods explicitly used this information. Consequently, these methods recovered target genes of deleted transcription factors more reliably than the inference methods that did not leverage this information (Fig. 2c). Explicit use of such knockouts also helped methods to more reliably draw the direction of edges between transcription factors. These observations suggest that measurements of transcription factor knockouts can be very informative for network reconstruction. In particular, this is the case for the *E. coli* dataset, which contained the largest number of such experiments (see **Methods**). To further explore the information content of different experiments, we employed a machine learning framework²² to systematically analyze the information gain from microarrays grouped according to the type of experimental perturbation (knockouts, drug perturbations, environmental perturbations, and time series; Supplementary Note 5). We found that experimental conditions independent of transcription factor knockout and overexpression also provide information, though at a reduced level.

Community networks outperform individual inference methods

Network inference methods have complementary advantages and limitations under different contexts, which suggests that combining the results of multiple inference methods could be a good strategy for improving predictions. We therefore integrated the predictions of all participating teams to construct community networks by re-scoring interactions according to their average rank across all methods (Supplementary Note 6). The integrated community network ranks 1st for *in silico*, 3rd for *E. coli*, and 6th for *S. cerevisiae* out of the 35 applied inference methods, which shows that the community network is consistently as good or better than the top individual methods (Fig. 2a). Thus, it has by far the best performance reflected in the overall score. We stress that, even though top-performing methods for a given network are competitive with the integrated community method, the performance of individual methods does not generalize across networks. Given the biological variation between organisms and the experimental variation between gene expression datasets, it is difficult to determine beforehand which methods will perform optimally for reconstructing an unknown regulatory network. In contrast, the community approach performs robustly across diverse datasets.

We next analyzed how the number of integrated methods affects the performance of community predictions by examining randomly sampled combinations of individual methods. On average, community methods perform better than individual inference methods even when integrating small sets of individual predictions, e.g., just five teams (Fig. 3a). Performance increases further with the number of integrated methods. For instance, given twenty inference methods, their integration ranks first or second 98% of the cases (Fig. 3b). We also found that the performance of the community network can be improved by increasing the diversity of the underlying inference methods. Consensus predictions from teams utilizing similar methodologies were outperformed by consensus predictions from diverse methodologies (Fig. 3c).

A key feature in taking a community network approach is robustness to the inclusion of a limited subset (up to ~20%) of poorly performing inference methods (Fig. 3d). Poor predictors essentially contributed noise, but this did not affect the performance of the community approach as a whole. This finding is crucial because the performance of individual methods when inferring regulatory networks for poorly studied organisms is not known *a priori* and is hard to evaluate empirically — even top performers on a benchmark network (e.g. *E. coli*) have varied performance when inferring a new, unknown network (e.g. *S. aureus*). On the other hand, adding good performers substantially increased the

performance of the community approach (Fig. 3d), which highlights the importance of developing high quality individual inference methods.

***E. coli* and *S. aureus* community networks**

To gain insights into transcriptional gene regulation for two bacteria, *E. coli* and *S. aureus*, we constructed networks for both organisms by integrating the predictions of all teams using the average rank method. Figure 4 shows the community networks for both organisms at a cutoff of 1,688 edges, which corresponds to an estimated precision of 50% for the *E. coli* network based on the gold standard of experimentally validated interactions from RegulonDB (**Methods**). At this cutoff, 50% of the *de novo* predicted regulatory edges were recovered known interactions; the remaining 50% may be false positives or newly discovered true interactions.

The precision of the *S. aureus* network cannot be measured accurately because there are comparatively few experimentally supported interactions available. Nevertheless, we confirmed the robustness of the consensus predictions by evaluating the network using the largely computationally-derived interactions from the RegPrecise database²³ (Supplementary Note 7).

We found that the *E. coli* and *S. aureus* networks both have a modular structure²⁴; that is, they comprise clusters of genes that are more densely connected amongst themselves than with other parts of the network. After identifying these modules²⁴, we tested them for enrichment of Gene Ontology terms (Supplementary Note 7). Network modules are strongly enriched for very specific biological processes. This allowed us to assign unique functions to most of the identified modules in both networks (Fig. 4 and Supplementary Data 6). As a specific example of an enriched module, 27 genes in *S. aureus* are highly enriched for pathogenic genes (Fig. 4b). These include exotoxins (*set7*, *set8*, *set11*, *set14*), genes responsible for biofilm formation (*tcaR*) and antibiotic metabolism (*tetR*), as well as a cell surface protein (*fnb*). The remaining 20 genes of this module are uncharacterized, but the predicted connections suggest their role in pathogenesis. This example illustrates how the inferred networks generate specific hypotheses regarding both the regulation and function of uncharacterized genes, enabling targeted validation efforts.

Experimental support of novel interactions

In addition to validation against known interactions from the RegulonDB gold standard, we experimentally tested a subset of novel predictions from the *E. coli* community network described above. We selected 5 transcription factors (*rhaR*, *cueR*, *purR*, *mprA*, and *gadE*), and then individually tested each of the 53 corresponding target gene predictions (Supplementary Note 8). Using qPCR, we measured the expression of each predicted target gene in the absence and presence of a chemical inducer known to activate the corresponding transcription factor (rhamnose for *rhaR*, copper sulfate for *cueR*, adenine for *purR*, carbonyl cyanide *m*-chlorophenylhydrazone for *mprA*, and hydrogen chloride for *gadE*). To control for possible indirect transcriptional responses, we also measured target gene expression in transcription factor deletion strains, again in the absence and presence of the chemical inducer. Putative targets were considered confirmed if they showed (1) strong response to the inducer of the respective transcription factor in the wild type and (2) no response to the inducer in the transcription factor deletion strain. We observed a clear difference between the two responses (>1.8 fold) for 23 novel targets out of 53 tested (Fig. 4c); this corresponds to a precision of ~40% for novel interactions, which is in line with our estimate of ~50% precision based on known interactions from RegulonDB. We note that these data support a direct regulatory effect of the tested transcription factor on the target gene, but chromatin immunoprecipitation experiments would be required to determine physical binding.

We observe a large variation in experimental validation among individual transcription factors (Fig. 4c). For *purR*, a key regulator in purine nucleotide metabolism, 10 of the 12 predicted target genes were experimentally supported. Nucleotide metabolism is a fundamental biological process that is affected across multiple conditions, thus *purR* regulation is well sampled across the *E. coli* dataset. However, in the case of *rhaR*, a key regulator in L-rhamnose degradation, none of the novel target gene predictions showed signs of regulation. L-rhamnose degradation is a specialized process that is only activated in the presence of L-rhamnose, and there were no conditions in the *E. coli* dataset where L-rhamnose degradation was explicitly tested. In the instance of *cueR*, a transcriptional regulator activated in the presence of copper, 4 out of 7 novel target gene predictions were confirmed. As with *rhaR*, there were no conditions in the dataset that explicitly tested copper regulation, yet unlike *rhaR*, network inference methods were able to identify true positive *cueR* regulatory interactions. These results suggest that while the overall precision for the network is high, the reliability of predictions for individual transcription factors can vary. When constructing a compendium of microarrays for global network inference, biases towards oversampling a narrow set of experimental conditions should thus be avoided.

Discussion

The DREAM project provides a unique framework where network inference methods from a community of experts are collected and impartially assessed on benchmark datasets. The collection of 35 inference methods assessed here itself constitutes a unique resource, as it spans all commonly used approaches in the field. In addition, the collection includes novel approaches (including the two best individual team performers of the challenge), representing a snapshot of the latest developments in the field.

Our analyses revealed specific advantages and limitations of different inference approaches (see Supplementary Note 9 and the full description of approaches in Supplementary Note 10). Sparse linear regression methods performed well, but only when data resampling strategies such as bootstrapping were used (the best performing regression methods all used data resampling, while the worst performing methods did not). Sparsity constraints employed by these methods effectively increased performance for cascade motifs, at the cost of missing interactions in feed-forward loops, fan-in, and fan-out motifs. Bayesian network methods exhibited below-average performance in this challenge, likely because they use heuristic searches, which are often too costly for systematic data resampling and may be better suited for smaller networks. Information theoretic methods performed better than correlation-based methods, but the two approaches had similar biases in predicting regulatory relationships. Compared to regression and Bayesian network methods, they perform better on feed-forward loops, fan-ins, and fan-outs (the more densely connected parts of the network), but have an increased rate of false positives for cascades. Meta predictors performed more robustly across datasets than other categories of methods, however, they could not match the robustness and performance of the community predictions, likely because they combine methods that do not provide sufficient diversity. Among all categories, methods that made explicit use of direct transcription factor perturbations (knockout or overexpression) greatly improved prediction accuracy for downstream targets (albeit at an increased false positive rate for cascades). For improving individual inference approaches we suggest the following: (1) optimally exploit direct transcription factor perturbations; (2) employ strategies to avoid over-fitting, such as data resampling; (3) develop more effective approaches to distinguish direct from indirect regulation (feed-forward loops *vs.* cascades).

Overall, methods performed well for the *in silico* and prokaryotic (*E. coli*) datasets; however, inferring gene regulatory networks from the eukaryotic (*S. cerevisiae*) dataset

proved to be a greater challenge. A fundamental assumption of network inference algorithms is that mRNA levels of transcription factors and their targets tend to be correlated — we found that this is true for *E. coli*, but not for *S. cerevisiae* (Supplementary Note 5). While the lower coverage of *S. cerevisiae* gold standards may also play a role (*E. coli* has the best-known regulatory network of any free-living organism¹⁶), the poor correlation at the mRNA level in *S. cerevisiae* is likely due to the increased regulatory complexity and prevalence of post-transcriptional regulation in eukaryotes, suggesting that accurate inference of eukaryotic regulatory networks requires additional inputs, such as promoter sequences, transcription factor binding, and chromatin modification datasets⁷.

Individual studies that introduce a novel inference method naturally tend to focus on its advantages in a particular application, which can paint an over-optimistic picture of performance¹³. While previous studies have explored strengths and weaknesses of inference approaches^{2,3}, the present assessment further shows that method performance is not robust across species and varies greatly even in the same category of inference methods (Table 1). This implies that performance is more related to the details of implementation, rather than the choice of the underlying methodology.

In network inference, variation in performance presents a problem, but at the same time offers a solution. By integrating the predictions from individual methods into community networks, we show that advantages of different methods complement each other, while limitations tend to be cancelled out. Instead of relying on a single inference method with uncertain performance on a previously unseen network, integrating predictions across inference methods becomes the best strategy. We note that not all of the 29 methods are required for enhanced performance. By considering complementary methods, we have shown that performance can be significantly improved with as few as three methods (Fig. 3c).

Ensemble-based methods have a storied past, with applications ranging from economics¹ to machine learning²⁵. In systems biology, robust models are often constructed from ensembles of instances (e.g., different parameterizations or model structures) that are derived from experimental data via a single approach^{26–30}, such as Monte Carlo sampling. In contrast, we formed consensus predictions from a large array of heterogeneous inference approaches. These “meta predictors” have been successful in other machine learning competitions^{31,32}. We have observed from previous DREAM challenges anecdotal evidence that community predictions can rank amongst the top performers¹³, but we did not previously attempt a systematic study of prediction integration for network inference. Here we established, through rigorous assessments and experimentally derived datasets, the performance robustness of prediction integration for transcriptional gene network inference.

The shortcomings of individual methods revealed in our assessment present many opportunities for improving these methods. We also expect further improvements in performance from advanced community approaches that: (i) actively leverage the method-specific advantages with regard to the datasets and networks of interest; (ii) optimize diversity in the ensemble, e.g., by weighting methods so as to balance the contribution of different method categories or PCA clusters; and (iii) employ more sophisticated voting schemes to negotiate consensus networks. To help spur developments in these areas, we provide the GP-DREAM web platform for the community to develop and apply network inference and consensus methods (<http://dream.broadinstitute.org>). We will continue to expand this free toolkit with top performing methods from the DREAM challenges, as well as other methods contributed by the community.

Methods

Expression data and gold standards

The design of the DREAM5 network inference challenge is outlined in Figure 1 (full description in Supplementary Note 1). Affymetrix gene expression datasets were compiled for *E. coli*, *S. aureus*, and *S. cerevisiae* from the Gene Expression Omnibus (GEO) database⁵⁰. Microarray datasets were uniformly normalized using Robust Multichip Averaging (RMA)⁵¹. Each dataset queries the underlying regulatory network in hundreds of different conditions, ranging from time courses to gene, drug, and environmental perturbations. Note that the number of measurements of transcription factor specific perturbations varies among the datasets (*S. aureus*: 0/161, *E. coli*: 67/806 and yeast: 3/537). The fourth dataset is an *in silico* counterpart to the *E. coli* dataset, generated using GeneNetWeaver^{52,53} (version 4.0). The structure of the *in silico* network corresponds to the *E. coli* transcriptional regulatory network from RegulonDB¹⁶ (10% random edges were added, resulting in 3,940 interactions). In addition to the gene expression data, we provide a list of putative transcription factors for each dataset and a number of descriptive features for each microarray experiment (e.g., the target of a gene deletion, or the time point of a time-series experiment). It is important to note that the identity of the organisms from which the data was generated was unknown to the participants. This was achieved by encrypting certain aspects of the data, and by anonymizing gene names.

Participants were presented the challenge to infer direct regulatory interactions between transcription factors and target genes from the given gene expression datasets. The submission format was a ranked list of predicted regulatory relationships for each network³.

The gold standard set of known transcriptional interactions for *E. coli* was obtained from RegulonDB¹⁶. We only included well-established interactions annotated with “strong evidence” according to RegulonDB evidence classification (2,066 interactions). For *S. cerevisiae*, we considered several alternative gold standards derived from orthogonal datasets, namely CHIP binding data and evolutionary conserved transcription factor binding motifs¹⁸, as well as systematic transcription factor deletions⁵⁴ (Supplementary Note 3). For the results reported in the main text, we used the most stringent gold standard, which includes only interactions that have both strong evidence of binding and conservation¹⁸.

All data and scripts are available in Supplementary Data 1 and at the DREAM website: <http://wiki.c2b2.columbia.edu/dream/index.php/D5c4>. The original microarray datasets are also publically available at the Many Microbe Microarrays Database⁵⁵ (M3D, <http://m3d.bu.edu/dream>).

Performance metrics

A detailed description of all performance metrics is given in Supplementary Note 4. Briefly, transcription factor-target predictions were evaluated as a binary classification task. The gold-standard networks represent the true positive interactions; the remaining pairs are considered negatives. Only the top 100,000 edge predictions were accepted. Pairs of nodes not part of the submitted list were considered to appear randomly ordered at the end of the list. Performance was assessed using the area under the ROC curve (AUROC) and the area under the precision vs. recall curve (AUPR)¹⁴. Note that predictions for genes that are not part of the gold standard, i.e., for which no experimentally supported interactions exist, were ignored in this evaluation.

AUROC and AUPR were separately transformed into *p*-values by simulating a null distribution for 25,000 random networks. Random edge lists were constructed by sampling edges from the submitted edge lists of the participants and assigning these edges random

ranks between 1 and 100,000. The histogram of randomly obtained AUROC and AUPR values was fit using stretched exponentials to extrapolate the distribution to values beyond the immediate range of the histogram¹⁴. To compute an overall score that summarizes the performance over the three networks with available gold standards (*E. coli*, *S. cerevisiae* and *in silico*), we used the same metric as in the previous two editions of the challenge^{3,14}, which is defined as the mean of the (log-transformed) network specific p -values:

$$\begin{aligned} \text{score}_{\text{ROC}} &= \frac{1}{3} \sum_{i=1}^3 -\log_{10} p_{\text{ROC},i} \\ \text{score}_{\text{PR}} &= \frac{1}{3} \sum_{i=1}^3 -\log_{10} p_{\text{PR},i} \\ \text{score} &= \frac{\text{score}_{\text{ROC}} + \text{score}_{\text{PR}}}{2} \end{aligned}$$

Clustering of inference approaches by principal component analysis (PCA)

We constructed a prediction matrix P , where rows correspond to edges (transcription factor-target pairs) and columns to inference methods. The element $p_{i,j}$ of this matrix is thus the rank assigned to edge i by inference method j . We only considered edges that figured in the top 100,000 predicted edges of at least three inference methods, yielding 1,175,525 interactions across the four datasets. Note that knowledge of a gold standard network is not required for the PCA, thus the *S. aureus* predictions were included in this analysis. The dimensionality of the combined prediction matrix (including the predictions for all four datasets) was reduced by PCA using SVDLIBC with standard parameters (<http://tedlab.mit.edu/~dr/SVDLIBC>). Results are consistent when performing PCA for each of the four datasets separately (Supplementary Note 4).

Network motif analysis

The goal of the network motif analysis is to evaluate, for a given network inference method, whether some types of edges of motifs are systematically predicted less (or more) reliably than expected³. We considered the six motif types illustrated in Figure 2. For each type of motif m , we identified all instances in the gold standard network and determined the average rank $r_{\bar{m}}$ assigned to its edges by the inference method. We further determined the average rank assigned to all edges that are *not* part of this motif type. The prediction bias is given by the difference $r_m - r_{\bar{m}}$. See Supplementary Note 4 for details.

Experimental materials and design

Novel predictions were selected from the *E. coli* community network with greater than 50% predicted precision. Transcription factors with at least 8 novel predictions were selected, including rhaR, cueR, purR, mprA, and gadE (note that the dataset supplied to the DREAM5 participants did not contain any knockout measurement for these transcription factors). Primers were designed for all novel target gene predictions after accounting for operon structure and at least 1 known target of the transcription factor was included as a positive control. A total of 53 predictions and 6 positive controls were tested.

For each transcription factor, a knockout strain was generated from the background *E. coli* strain BW25113. Each transcription factor was induced by a different stimulus: rhamnose for rhaR, copper sulfate for cueR, adenine for purR, carbonyl cyanide *m*-chlorophenylhydrazone for mprA, and HCl for gadE. Four experimental conditions were used for each transcription factor: background strain without inducer (WT(-)), background strain with inducer (WT(+)), deletion strain without inducer (Δ (-)), and deletion strain with inducer (Δ (+)). Three biological replicates were generated for all experimental conditions. Cultures were grown in LB media or minimal media (Supplementary Note 8), and

incubation was performed in darkened shakers (300 RPM) at 37°C. PCR primers were designed for all target genes. Target genes were quantified through qPCR using LightCycler 480 SYBR Green I Master Kit (Roche Applied Science). True positive interactions were expected to meet two criteria: (1) a strong response to the TF inducer in wild type, and (2) no or weak response to the TF inducer in the TF-deletion strain. Target gene interactions were considered to have “strong support” if the ratio of criteria 1 to criteria 2, $(WT(+)/WT(-)) / (\Delta(+)/\Delta(-))$, was greater than two and “weak support” if the ratio was greater than 1.8 (Supplementary Data 7).

Supplementary Material

Refer to Web version on PubMed Central for supplementary material.

Acknowledgments

We thank all challenge participants for their invaluable contribution. We thank R. Norel and J. Saez Rodriguez who participated in different aspects of the organization and scoring of DREAM5. We also thank P. Carr, M. Reich, J. Mesirov, and the rest of the GenePattern team for providing software and support. This work was in part funded by the National Institutes of Health National Centers for Biomedical Computing Roadmap Initiative (U54CA121852), Howard Hughes Medical Institute, National Institutes of Health Director’s Pioneer Award DPI OD003644, and a fellowship from the Swiss National Science Foundation to D.M. Challenge participants acknowledge: Grants no. ANR-07-BLAN-0311-03 and ANR-09-BLAN-0051-04 from the French National Research Agency (A.-C.H., P.V.-L., F.M., J.-P.V.); the Interuniversity Attraction Poles Programme (IAP P6/25 BIOMAGNET), initiated by the Belgian State, Science Policy Office, the French Community of Belgium (ARC Biomod), and the European Network of Excellence PASCAL2 (V.A.H.-T., A.I., L.W., Y.S., P.G.); V.A.H.-T. is recipient of a fellowship from the Fonds pour la formation à la Recherche dans l’Industrie et dans l’Agriculture (F.R.I.A., Belgium); Y.S. is a postdoctoral fellow of the Fonds voor Wetenschappelijk Onderzoek -Vlaanderen (FWO, Belgium); P.G. is Research Associate of the Fonds National de la Recherche Scientifique (FNRS, Belgium); the European Community’s 7th Framework Program, grant no. HEALTH-F4-2007-200767 for the APO-SYS program, and a doctoral fellowship from the Edmond J. Safra Bioinformatics Program at Tel Aviv University (G.K., R.S.); the Irish Research Council for Science Engineering and Technology for financial support under the EMBARK scheme, and the Irish Centre for High-End Computing for provision of computational facilities and technical support (A. Sirbu, H.J.R., M.C.); the US National Cancer Institute grant no. U54CA132383 and US National Science Foundation grant no. HRD-0420407 (Z.O., Y.Z., H.W., M.S.); and the Sardinian Regional Authorities (A.F., A.P., N.S., V.L.).

References

1. Surowiecki, J. *The Wisdom of Crowds : Why the Many are Smarter than the Few and How Collective Wisdom Shapes Business, Economies, Societies, and Nations.* 2004.
2. De Smet R, Marchal K. Advantages and limitations of current network inference methods. *Nat Rev Microbiol.* 2010; 8:717–729. [PubMed: 20805835]
3. Marbach D, et al. Revealing strengths and weaknesses of methods for gene network inference. *Proc Natl Acad Sci USA.* 2010; 107:6286–6291. [PubMed: 20308593]
4. Bar-Joseph Z, et al. Computational discovery of gene modules and regulatory networks. *Nat Biotechnol.* 2003; 21:1337–1342. [PubMed: 14555958]
5. Reiss DJ, Baliga NS, Bonneau R. Integrated biclustering of heterogeneous genome-wide datasets for the inference of global regulatory networks. *BMC Bioinformatics.* 2006; 7:280. [PubMed: 16749936]
6. Lemmens K, et al. DISTILLER: a data integration framework to reveal condition dependency of complex regulons in *Escherichia coli*. *Genome Biol.* 2009; 10:R27. [PubMed: 19265557]
7. Marbach D, et al. Predictive regulatory models in *Drosophila melanogaster* by integrative inference of transcriptional networks. *Genome research.* 2012; 22:1101/gr.127191.111
8. Friedman N, Linial M, Nachman I, Pe’er D. Using Bayesian networks to analyze expression data. *J Comput Biol.* 2000; 7:601–620. [PubMed: 11108481]
9. Margolin AA, et al. ARACNE: an algorithm for the reconstruction of gene regulatory networks in a mammalian cellular context. *BMC Bioinformatics.* 2006; 7 (Suppl 1):S7. [PubMed: 16723010]

10. di Bernardo D, et al. Chemogenomic profiling on a genome-wide scale using reverse-engineered gene networks. *Nat Biotechnol.* 2005; 23:377–383. [PubMed: 15765094]
11. Faith JJ, et al. Large-scale mapping and validation of *Escherichia coli* transcriptional regulation from a compendium of expression profiles. *PLoS Biol.* 2007; 5:e8. [PubMed: 17214507]
12. Stolovitzky G, Monroe D, Califano A. Dialogue on reverse-engineering assessment and methods: the DREAM of high-throughput pathway inference. *Ann N Y Acad Sci.* 2007; 1115:1–22. [PubMed: 17925349]
13. Stolovitzky G, Prill RJ, Califano A. Lessons from the DREAM2 Challenges. *Ann N Y Acad Sci.* 2009; 1158:159–195. [PubMed: 19348640]
14. Prill RJ, et al. Towards a rigorous assessment of systems biology models: the DREAM3 challenges. *PLoS ONE.* 2010; 5:e9202. [PubMed: 20186320]
15. Reich M, et al. GenePattern 2.0. *Nat Genet.* 2006; 38:500–501. [PubMed: 16642009]
16. Gama-Castro S, et al. RegulonDB version 7.0: transcriptional regulation of *Escherichia coli* K-12 integrated within genetic sensory response units (Gensor Units). *Nucleic Acids Res.* 2011; 39:D98–105. [PubMed: 21051347]
17. Harbison CT, et al. Transcriptional regulatory code of a eukaryotic genome. *Nature.* 2004; 431:99–104. [PubMed: 15343339]
18. MacIsaac KD, et al. An improved map of conserved regulatory sites for *Saccharomyces cerevisiae*. *BMC Bioinformatics.* 2006; 7:113. [PubMed: 16522208]
19. Huynh-Thu VA, Irrthum A, Wehenkel L, Geurts P. Inferring regulatory networks from expression data using tree-based methods. *PLoS ONE.* 2010; 5
20. Küffner R, Petri T, Tavakkolkhah P, Windhager L, Zimmer R. Inferring Gene Regulatory Networks by ANOVA. *Bioinformatics.* 2012; 28:1376–1382. [PubMed: 22467911]
21. Tibshirani R. Regression shrinkage and selection via the Lasso. *J R Stat Soc Ser B-Methodol.* 1996; 58:267–288.
22. Mordelet F, Vert JP. SIRENE: supervised inference of regulatory networks. *Bioinformatics.* 2008; 24:i76–82. [PubMed: 18689844]
23. Ravcheev DA, et al. Inference of the Transcriptional Regulatory Network in *Staphylococcus aureus* by Integration of Experimental and Genomics-Based Evidence. *Journal of Bacteriology.* 2011; 193:3228–3240. [PubMed: 21531804]
24. Newman MEJ. Modularity and community structure in networks. *Proc Natl Acad Sci USA.* 2006; 103:8577–8582. [PubMed: 16723398]
25. Dietterich TG. Ensemble Methods in Machine Learning. *Multiple Classifier Systems.* 2000; 1857:1–15.
26. Prinz AA, Bucher D, Marder E. Similar network activity from disparate circuit parameters. *Nat Neurosci.* 2004; 7:1345–1352. [PubMed: 15558066]
27. Kuepfer L, Peter M, Sauer U, Stelling J. Ensemble modeling for analysis of cell signaling dynamics. *Nat Biotechnol.* 2007; 25:1001–1006. [PubMed: 17846631]
28. Kaltenbach HM, Dimopoulos S, Stelling J. Systems analysis of cellular networks under uncertainty. *FEBS Lett.* 2009; 583:3923–3930. [PubMed: 19879267]
29. Marbach D, Mattiussi C, Floreano D. Combining multiple results of a reverse-engineering algorithm: application to the DREAM five-gene network challenge. *Ann N Y Acad Sci.* 2009; 1158:102–113. [PubMed: 19348636]
30. Marder E, Taylor AL. Multiple models to capture the variability in biological neurons and networks. *Nat Neurosci.* 2011; 14:133–138. [PubMed: 21270780]
31. Moult J. A decade of CASP: progress, bottlenecks and prognosis in protein structure prediction. *Current Opinion in Structural Biology.* 2005; 15:285–289. [PubMed: 15939584]
32. Bell RM, Koren Y. Lessons from the Netflix Prize Challenge. *SIGKDD Explorations.* 2007; 9:75–79.
33. Haury, A-C.; Mordelet, F.; Vera-Licona, P.; Vert, J-P. TIGRESS: Trustful Inference of Gene REgulation using Stability Selection. 2012. arXiv:1205.1181at <<http://arxiv.org/abs/1205.1181>>
34. Yuan M, Lin Y. Model selection and estimation in regression with grouped variables. *J Royal Statist Soc B.* 2006; 68:49–67.

35. Lèbre S, Becq J, Devaux F, Stumpf MPH, Lelandais G. Statistical inference of the time-varying structure of gene-regulation networks. *BMC Syst Biol.* 2010; 4:130. [PubMed: 20860793]
36. Meinshausen N, Bühlmann P. Stability selection. *J Royal Statist Soc B.* 2010; 72:417–473.
37. van Someren EP, et al. Least absolute regression network analysis of the murine osteoblast differentiation network. *Bioinformatics.* 2006; 22:477–484. [PubMed: 16332709]
38. Butte AJ, Kohane IS. Mutual information relevance networks: functional genomic clustering using pairwise entropy measurements. *Pac Symp Biocomput.* 2000:418–429. [PubMed: 10902190]
39. Mani S, Cooper GF. A Bayesian Local Causal Discovery Algorithm. *Proceedings of the World Congress on Medical Informatics.* 2004:731–735.
40. Tsamardinos I, Aliferis CF, Statnikov A. Time and sample efficient discovery of Markov blankets and direct causal relations. *Proceedings of the ninth ACM SIGKDD international conference on Knowledge discovery and data mining.* 2003:673–678.10.1145/956750.956838
41. Aliferis CF, Statnikov A, Tsamardinos I, Mani S, Koutsoukos XD. Local causal and Markov blanket induction for causal discovery and feature selection for classification. Part I: Algorithm and Empirical Evaluation. *Journal of Machine Learning Research.* 2010; 11:171–234.
42. Statnikov A, Aliferis CF. Analysis and computational dissection of molecular signature multiplicity. *PLoS Comput Biol.* 2010; 6:e1000790. [PubMed: 20502670]
43. Karlebach G, Shamir R. Constructing logical models of gene regulatory networks by integrating transcription factor-DNA interactions with expression data: an entropy-based approach. *J Comput Biol.* 2012; 19:30–41. [PubMed: 22216865]
44. Yeung KY, Bumgarner RE, Raftery AE. Bayesian model averaging: development of an improved multi-class, gene selection and classification tool for microarray data. *Bioinformatics.* 2005; 21:2394–2402. [PubMed: 15713736]
45. Yip KY, Alexander RP, Yan KK, Gerstein M. Improved reconstruction of *in silico* gene regulatory networks by integrating knockout and perturbation data. *PLoS ONE.* 2010; 5:e8121. [PubMed: 20126643]
46. Sîrbu A, Ruskin HJ, Crane M. Stages of Gene Regulatory Network Inference: the Evolutionary Algorithm Role. *Evolutionary Algorithms.* 2011
47. Song MJ, et al. Reconstructing generalized logical networks of transcriptional regulation in mouse brain from temporal gene expression data. *EURASIP J Bioinform Syst Biol.* 2009:545176.10.1155/2009/545176 [PubMed: 19300527]
48. Greenfield A, Madar A, Ostrer H, Bonneau R. DREAM4: Combining genetic and dynamic information to identify biological networks and dynamical models. *PLoS ONE.* 2010; 5:e13397. [PubMed: 21049040]
49. Watkinson J, Liang KC, Wang X, Zheng T, Anastassiou D. Inference of regulatory gene interactions from expression data using three-way mutual information. *Ann N Y Acad Sci.* 2009; 1158:302–313. [PubMed: 19348651]
50. Barrett T, et al. NCBI GEO: archive for functional genomics data sets--10 years on. *Nucleic Acids Res.* 2011; 39:D1005–1010. [PubMed: 21097893]
51. Bolstad BM, Irizarry RA, Astrand M, Speed TP. A comparison of normalization methods for high density oligonucleotide array data based on variance and bias. *Bioinformatics.* 2003; 19:185–193. [PubMed: 12538238]
52. Marbach D, Schaffter T, Mattiussi C, Floreano D. Generating realistic *in silico* gene networks for performance assessment of reverse engineering methods. *J Comput Biol.* 2009; 16:229–239. [PubMed: 19183003]
53. Schaffter T, Marbach D, Floreano D. GeneNetWeaver: *in silico* benchmark generation and performance profiling of network inference methods. *Bioinformatics.* 2011; 27:2263–2270. [PubMed: 21697125]
54. Hu Z, Killion PJ, Iyer VR. Genetic reconstruction of a functional transcriptional regulatory network. *Nat Genet.* 2007; 39:683–687. [PubMed: 17417638]
55. Faith JJ, et al. Many Microbe Microarrays Database: uniformly normalized Affymetrix compendia with structured experimental metadata. *Nucl Acids Res.* 2008; 36:D866–870. [PubMed: 17932051]

The DREAM5 consortium

Andrej Aderhold^{1, 2}, Kyle R. Allison³, Richard Bonneau^{4, 5, 6}, Diogo M. Camacho³, Yukun Chen⁷, James J. Collins³, Francesca Cordero^{8, 9}, James C. Costello³, Martin Crane¹⁰, Frank Dondelinger^{1, 11}, Mathias Drton¹², Roberto Esposito⁸, Rina Foygel¹², Alberto de la Fuente¹³, Jan Gertheiss¹⁴, Pierre Geurts^{15, 16}, Alex Greenfield⁵, Marco Grzegorzczak¹⁷, Anne-Claire Haury^{18, 19, 20}, Benjamin Holmes^{21, 22}, Torsten Hothorn¹⁴, Dirk Husmeier¹, Vân Anh Huynh-Thu^{15, 16}, Alexandre Irrthum^{15, 16}, Manolis Kellis^{21, 22}, Guy Karlebach²³, Robert Küffner²⁴, Sophie Lèbre²⁵, Vincenzo De Leo^{13, 26}, Aviv Madar⁴, Subramani Mani⁷, Daniel Marbach^{21, 22}, Fantine Mordelet^{18, 19, 20, 27}, Harry Ostrer²⁸, Zhengyu Ouyang²⁹, Ravi Pandya³⁰, Tobias Petri²⁴, Andrea Pinna¹³, Christopher S. Poultnery⁴, Robert J. Prill³¹, Serena Rezny¹², Heather J. Ruskin¹⁰, Yvan Saeys^{32, 33}, Ron Shamir²³, Alina Sîrbu¹⁰, Mingzhou Song²⁹, Nicola Soranzo¹³, Alexander Statnikov³⁴, Gustavo Stolovitzky³¹, Nicci Vega³, Paola Vera-Licona^{18, 19, 20}, Jean-Philippe Vert^{18, 19, 20}, Alessia Visconti⁸, Haizhou Wang²⁹, Louis Wehenkel^{15, 16}, Lukas Windhager²⁴, Yang Zhang²⁹, and Ralf Zimmer²⁴

¹Biomathematics and Statistics Scotland, Edinburgh & Aberdeen, UK

²School of Biology, University of St Andrews, UK

³Howard Hughes Medical Institute, Department of Biomedical Engineering, and Center for BioDynamics, Boston University, Boston, Massachusetts, USA.

⁴Department of Biology, Center for Genomics & Systems Biology, New York University, New York, NY, USA

⁵Computational Biology Program, New York University Sackler School of Medicine, New York, NY, USA

⁶Computer Science Department, Courant Institute of Mathematical Sciences, New York University, New York, NY, USA

⁷Department of Biomedical Informatics, Vanderbilt University, Nashville, TN, USA

⁸Department of Computer Science, Corso Svizzera 185, 10149, Torino, Italy

⁹Department of Clinical and Biological Sciences, Regione Gonzole 10, Orbassano, Italy

¹⁰Centre for Scientific Computing and Complex Systems Modelling, School of Computing, Dublin City University, Ireland

¹¹School of Informatics, University of Edinburgh, UK

¹²University of Chicago, Department of Statistics, Chicago, IL, USA

¹³CRS4 Bioinformatica, Parco Tecnologico della Sardegna, Edificio 3, Loc. Piscina Manna, 09010 Pula (CA), Italy

¹⁴Ludwig-Maximilians University, Department of Statistics, Ludwigstr. 33, 80539 Munich, Germany

¹⁵Department of Electrical Engineering and Computer Science, Systems and Modeling, University of Liège, Belgium

¹⁶GIGA-Research, Bioinformatics and Modeling, University of Liège, Belgium

¹⁷Department of Statistics, TU Dortmund University, Germany

¹⁸Mines ParisTech, CBIO, 35 rue Saint-Honoré, Fontainebleau, F-77300, France

¹⁹Institut Curie, Paris, F-75248, France

²⁰INSERM, U900, Paris, F-75248, France

²¹Computer Science and Artificial Intelligence Laboratory (CSAIL), Massachusetts Institute of Technology (MIT), Cambridge, Massachusetts, USA.

²²Broad Institute of MIT and Harvard, Cambridge, Massachusetts, USA.

²³The Blavatnik School of Computer Science, Tel Aviv University, Israel

²⁴Ludwig-Maximilians University, Department of Informatics, Amalienstr. 17, 80333 Munich, Germany

²⁵LSIIT, UMR UdS-CNRS 7005, Illkirch Université de Strasbourg, France

²⁶Linkalab, Complex Systems Computational Laboratory, 09100 Cagliari (CA), Italy

²⁷CREST, INSEE, Malakoff, F-92240, France

²⁸Human Genetics Program, Department of Pediatrics, New York University Langone Medical Center, New York, NY, USA

²⁹Department of Computer Science, New Mexico State University, Las Cruces, NM, USA

³⁰Microsoft Research, 1 Microsoft Way, Redmond, WA, USA

³¹IBM T. J. Watson Research Center, Yorktown Heights, New York, USA.

³²Department of Plant Systems Biology, VIB, Gent, Belgium

³³Department of Plant Biotechnology and Bioinformatics, Ghent University, Ghent, Belgium

³⁴Center for Health Informatics and Bioinformatics, New York University, New York, NY, USA

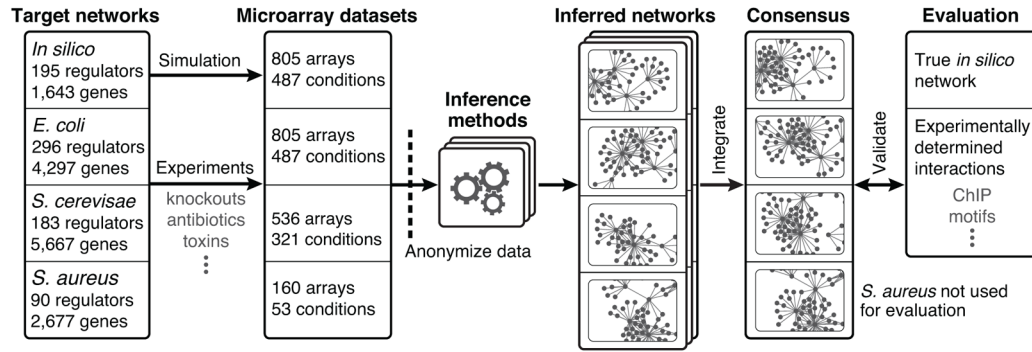


Figure 1. The DREAM5 network inference challenge

Assessment involved the following steps (from left to right). (1) Participants were challenged to infer the genome-wide transcriptional regulatory networks of *E. coli*, *S. cerevisiae*, and *S. aureus*, as well as an *in silico* (simulated) network. (2) Gene expression datasets for a wide range of experimental conditions were compiled. Anonymized datasets were released to the community, hiding the identities of the genes. (3) 29 participating teams inferred gene regulatory networks. In addition, we applied 6 “off-the-shelf” inference methods. (4) Network predictions from individual teams were integrated to form community networks. (5) Network predictions were assessed using experimentally supported interactions from *E. coli* and *S. cerevisiae*, as well as the known *in silico* network.

\$watermark-text

\$watermark-text

\$watermark-text

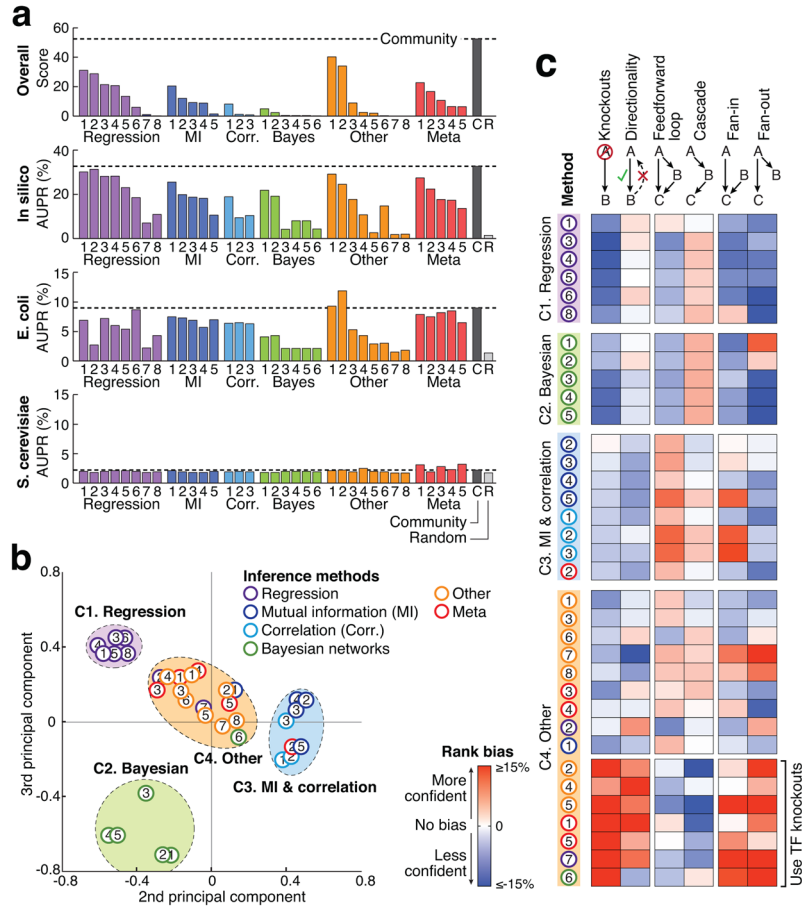


Figure 2. Evaluation of network inference methods

Inference methods are indexed according to Table 1. **(a)** The plots depict the performance for the individual networks (area under precision-recall curve, AUPR) and the overall score summarizing the performance across networks (**Methods**). R indicates performance of random predictions. C indicates performance of the integrated community predictions. **(b)** Methods are grouped according to the similarity of their predictions via principal component analysis. Shown are the 2nd vs. 3rd principal components; the 1st principal component accounts mainly for the overall performance (Supplementary Note 4). **(c)** The heatmap depicts method-specific biases in predicting network motifs. Rows represent individual methods and columns represent different types of regulatory motifs. Red and blue show interactions that are easier and harder to detect, respectively.




Figure 3. Analysis of community networks vs. individual inference methods

(a) The plot shows the overall score, which summarizes performance across the *E. coli*, *S. cerevisiae*, and *in silico* networks, for individual inference methods or various combinations of integrated methods. The first boxplot depicts the performance distribution of individual inference methods ($K=1$). Subsequent boxplots show the performance when integrating $K>1$ randomly sampled methods. The red bar shows the performance when integrating all methods ($K=29$). Boxplots depict performance distributions with respect to the minimum, the maximum and the three quartiles. (b) The probability that the community network ranks among the top $x\%$ of the K individual methods used to construct the community network. The diagonal shows the expected performance when choosing an individual method ($K=1$). (c) The integration of complementary methods is particularly beneficial. The first boxplot shows the performance of individual methods from clusters 1–3 (as defined in Fig. 2b). The second and third boxplots show performance of community networks obtained by integrating three randomly selected inference methods: (i) from the same cluster, or (ii) from different clusters. (d) The plots show the overall score for an initial community network formed by integrating all individual methods (open circles, blue) except for the best five and worst five. One-by-one the worst five (left panel) and best five (right panel) methods are added to form additional community networks (filled circles, red).

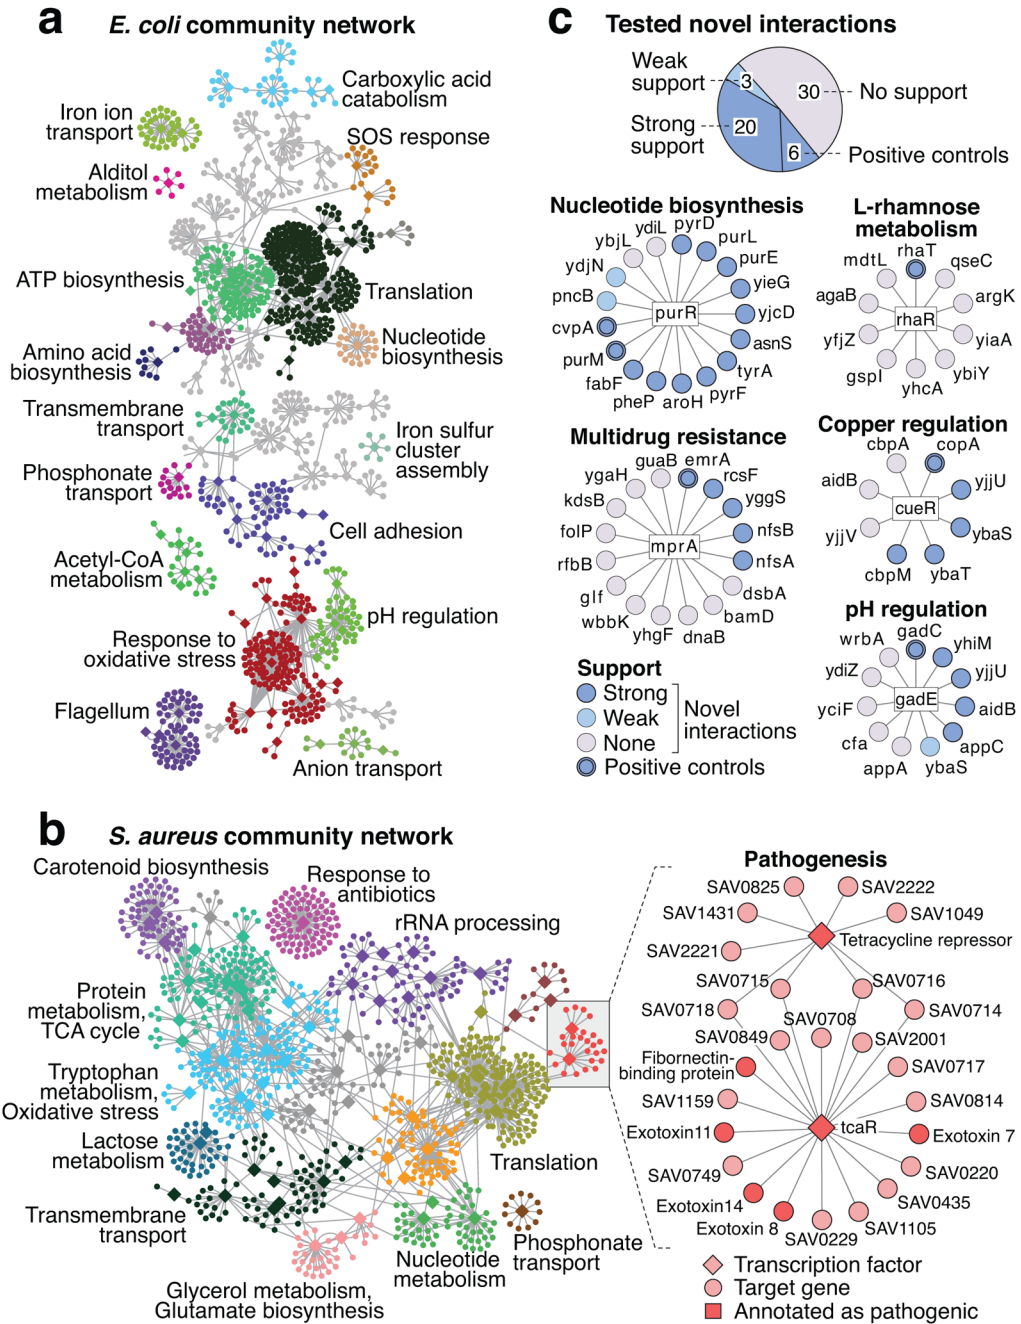


Figure 4. *E. coli* and *S. aureus* community networks

(a, b) At a cutoff of 1688 edges, the (a) *E. coli* community network connects 1,505 genes (including 204 transcription factors, shown as diamonds), and the (b) *S. aureus* network connects 1,084 genes (85 transcription factors). Network modules were identified and tested for Gene Ontology term enrichment, as indicated (grey colored genes do not show enrichment). A network module enriched for Gene Ontology terms related to pathogenesis is highlighted in the *S. aureus* network. (c) The schematics depict newly predicted *E. coli* regulatory interactions that were experimentally tested. The pie chart depicts the breakdown of strongly and weakly supported targets (Methods). The positive controls were six known interactions from RegulonDB.

Table 1

Network inference methods.

ID	Synopsis	Reference
Regression: Transcription factors are selected by target gene specific (1) sparse linear regression and (2) data resampling approaches.		
1	Trustful Inference of Gene REgulation using Stability Selection (TIGRESS): (1) Lasso; (2) the regularization parameter selects five transcription factors per target gene in each bootstrap sample.	33 ^a
2	(1) Steady state and time series data are combined by group lasso; (2) bootstrapping.	34 ^a
3	Combination of lasso and Bayesian linear regression models learned using Reversible Jump Markov Chain Monte Carlo simulations.	35 ^a
4	(1) Lasso; (2) bootstrapping.	36
5	(1) Lasso; (2) area under the stability selection curve.	36
6	Application of the Lasso toolbox GENLAB using standard parameters.	37
7	Lasso models are combined by the maximum regularization parameter selecting a given edge for the first time.	36 ^a
8	Linear regression determines the contribution of transcription factors to the expression of target genes.	— ^{a,b}
Mutual Information: Edges are (1) ranked based on variants of mutual information and (2) filtered for causal relationships.		
1	Context likelihood of relatedness (CLR): (1) Spline estimation of mutual information; (2) the likelihood of each mutual information score is computed based on its local network context.	11 ^{a,b}
2	(1) Mutual information is computed from discretized expression values.	38 ^{a,b}
3	Algorithm for the Reconstruction of Accurate Cellular Networks (ARACNE): (1) kernel estimation of mutual information; (2) the data processing inequality is used to identify direct interactions.	9 ^{a,b}
4	(1) Fast kernel-based estimation of mutual information; (2) Bayesian Local Causal Discovery (BLCD) and Markov blanket (HITON-PC) algorithm to identify direct interactions.	39 ^a
5	(1) Mutual information and Pearson's correlation are combined; (2) BLCD and HITON-PC algorithm.	39 ^a
Correlation: Edges are ranked based on variants of correlation.		
1	Absolute value of Pearson's correlation coefficient.	38
2	Signed value of Pearson's correlation coefficient.	38 ^{a,b}
3	Signed value of Spearman's correlation coefficient.	38 ^{a,b}
Bayesian networks optimize posterior probabilities by different heuristic searches.		
1	Simulated annealing (catnet R package, http://cran.r-project.org/web/packages/catnet), aggregation of three runs.	—
2	Simulated annealing (catnet R package, http://cran.r-project.org/web/packages/catnet).	—
3	Max-Min Parent and Children algorithm (MMPC), bootstrapped datasets.	40
4	Markov blanket algorithm (HITON-PC), bootstrapped datasets.	41
5	Markov boundary induction algorithm (TIE*), bootstrapped datasets.	42
6	Models transcription factor perturbation data and time series using dynamic Bayesian networks (Infer.NET toolbox, http://research.microsoft.com/infernet).	— ^a
Other Approaches: Network inference by heterogeneous and novel methods.		
1	Genie3: A random forest is trained to predict target gene expression. Putative transcription factors are selected as tree nodes if they consistently reduce the variance of the target.	19 ^a
2	Co-dependencies between transcription factors and target genes are detected by the non-linear correlation coefficient η^2 (two-way ANOVA). Transcription factor perturbation data are up-weighted.	20 ^a
3	Transcription factors are selected maximizing the conditional entropy for target genes, which are represented as Boolean vectors with probabilities to avoid discretization.	43 ^a

ID	Synopsis	Reference
4	Transcription factors are preselected from transcription factor perturbation data or by Pearson's correlation and then tested by iterative Bayesian Model Averaging (BMA).	44
5	A Gaussian noise model is used to estimate if the expression of a target gene changes in transcription factor perturbation measurements.	45
6	After scaling, target genes are clustered by Pearson's correlation. A neural network is trained (genetic algorithm) and parameterized (back-propagation).	46 ^a
7	Data is discretized by Gaussian mixture models and clustering (Ckmeans); Interactions are detected by generalized logical network modeling (χ^2 test).	47 ^a
8	The χ^2 test is applied to evaluate the probability of a shift in transcription factor and target gene expression in transcription factor perturbation experiments.	47 ^a
Meta predictors (1) apply multiple inference approaches and (2) compute aggregate scores.		
1	(1) Z-scores for target genes in transcription factor knockout data, time-lagged CLR for time series, and linear ordinary differential equation models constrained by lasso (Inferelator); (2) resampling approach.	48 ^a
2	(1) Pearson's correlation, mutual information, and CLR; (2) rank average.	—
3	(1) Calculates target gene responses in transcription factor knockout data, applies full-order, partial correlation and transcription factor-target co-deviation analysis; (2) weighted average with weights trained on simulated data.	— ^a
4	(1) CLR filtered by negative Pearson's correlation, least angle regression (LARS) of time series, and transcription factor perturbation data; (2) combination by z-scores.	49
5	(1) Pearson's correlation, differential expression (limma), and time series analysis (maSigPro); (2) Naïve Bayes.	— ^a

Methods have been manually categorized based on participant-supplied descriptions. Within each class, methods are sorted by overall performance (see Figure 2a). Note that generic references have been used if more specific ones were not available.

^aDetailed method description included in Supplementary Note 10;

^bOff-the-shelf algorithm applied by challenge organizers.