

A MULTI-FUNCTIONAL ANALYZER USES PARAMETER CONSTRAINTS TO IMPROVE THE EFFICIENCY OF MODEL-BASED GENE-SET ANALYSIS

BY ZHISHI WANG¹, QIULING HE², BRET LARGET³ AND
MICHAEL A. NEWTON⁴

University of Wisconsin, Madison

We develop a model-based methodology for integrating gene-set information with an experimentally-derived gene list. The methodology uses a previously reported sampling model, but takes advantage of natural constraints in the high-dimensional discrete parameter space in order to work from a more structured prior distribution than is currently available. We show how the natural constraints are expressed in terms of linear inequality constraints within a set of binary latent variables. Further, the currently available prior gives low probability to these constraints in complex systems, such as Gene Ontology (GO), thus reducing the efficiency of statistical inference. We develop two computational advances to enable posterior inference within the constrained parameter space: one using integer linear programming for optimization and one using a penalized Markov chain sampler. Numerical experiments demonstrate the utility of the new methodology for a multivariate integration of genomic data with GO or related information systems. Compared to available methods, the proposed multi-functional analyzer covers more reported genes without mis-covering nonreported genes, as demonstrated on genome-wide data from association studies of type 2 diabetes and from RNA interference studies of influenza.

1. Introduction. In statistical genomics, the gene list is a recurring data structure. We have in mind situations where experimental results amount to a collection of genes measured to have some property. Examples include the following: RNA expression studies, in which the property might be differential expression of the gene between two cell types; genome-wide RNA knock-down studies, in which the property is significant phenotypic alteration caused by RNA interference; chromatin studies recording genes in the vicinity of transcription factor binding sites or having certain epigenetic marks. In all cases, the reported gene list is really the result of inference from more basic experimental data. These more basic data may

Received October 2013; revised August 2014.

¹Supported in part by a fellowship from the Morgridge Institute of Research.

²Supported in part by a research assistantship from the National Institutes of Health (NIH) (R21 HG006568); currently employed by Novartis Pharmaceuticals.

³Supported in part by NIH R01 GM086887 and NSF DEB 0949121.

⁴Supported in part by NIH R21 HG006568.

Key words and phrases. Gene-set enrichment, Bayesian analysis, integer linear programming.

be available to support subsequent analyses, but we are concerned with the important and relatively common case in which the gene list itself is the primary data set brought forward for analysis.

The statistical question of central importance in the present paper is how to interpret the gene list in the context of preexisting biological knowledge about the functional properties of all genes, as these exogenous data are recorded in database systems, notably Gene Ontology (GO), the Kyoto Encyclopedia (KEGG) and the reactome, among others [The Gene Ontology Consortium (2000); Kanehisa and Goto (2000); Matthews et al. (2009); Gentleman et al. (2004)]. For us, exogenous data form a collection of gene sets, with each set equaling those genes previously determined, by some evidence, to have a specific biological property. Recently, for example, the full GO collection contained 16,527 sets (GO terms) annotating 17,959 human genes.⁵ Needless to say, genes are typically annotated to multiple gene sets (median 7 sets per gene among the genes annotated to sets which contain between 3 and 30 genes, e.g.), covering all sorts of functional properties. The task of *gene-set analysis* is to efficiently interpret the functional content of an experimentally-derived gene list by somehow integrating these endogenous and exogenous data sources [Khatri, Sirota and Butte (2012)].

Our starting point is an exciting development in the methodology of gene-set analysis. Model-based gene-set analysis (MGSA) expresses gene-level indicators of presence on the gene list as Bernoulli trials whose success probabilities are determined in a simple way by latent activity states of binary variables associated with the gene sets [Bauer, Gagneur and Robinson (2010); Bauer, Robinson and Gagneur (2011)]. Inference seeks to identify the *active* gene sets, as these represent functional drivers of the experimental data. Inference is computationally difficult because the activity state of a given gene set depends not only on experimental data for genes in that set, but also on the unknown activity states of all other gene sets that annotate these same genes. MGSA overcomes the problem through Bayesian inference implemented with an efficient Markov chain Monte Carlo (MCMC) sampler, and thus provides marginal posterior probabilities that each gene set is in the active state. The MGSA methodology is compelling. Because it treats all gene sets in the collection simultaneously, it provides a truly multivariate analysis of the exogenous data source, where most available approaches are univariate (one set at a time). Where set/set overlaps are a nuisance in most gene-set methodologies, MGSA utilizes them directly in modeling and inference. This accounts for pleiotropy, that genes have multiple biological functions, reduces the risk of spurious associations, and leads to cleaner output whereby a typical list of gene sets inferred to be active is simpler and exhibits less redundancy than in standard univariate analyses [Bauer, Gagneur and Robinson (2010); Newton, He and Kendzierski (2012)].

⁵Bioconductor package org.Hs.eg.db, version 2.8.0.

Our analysis reveals a feature of MGSA that adversely affects its statistical properties. In ever denser collections of gene sets, the MGSA prior distribution puts more and more mass on logically inconsistent joint activity states. As a result, data need to work ever harder to overcome this misguided prior probability. The effect is tangible; for a given amount of data, fewer truly activated gene sets are inferred to be active, compared to what is achievable with an alternative formulation. We propose a new methodology, the multi-functional analyzer (MFA), which aims to improve the statistical efficiency of MGSA. It uses two computational advances that enable posterior inference in the high-dimensional constrained space of joint activity states. One is an efficient MCMC sampling scheme constructed by penalizing the log-posterior in the unconstrained space and one is a discrete optimization scheme that translates the inference problem into an integer linear programming (ILP) problem.

We note that inference about gene-set activity states may be interesting from the general perspective of high-dimensional statistics. Typically, dependence among data from different inference units (sets, in this case) is considered a nuisance and testing aims to identify nonnull units (active sets) by some methodology that is robust to dependencies, since these dependencies are often difficult to estimate from available data. In the present context dependencies are complicated but explicit, and inference benefits by using them to advantage. Finally, we also note that the probability model underlying our methodology—the *role model*—has potential utility in other domains of application. It provides a simple way to relate data collected at one level (genes, in this case) to inference units that are unordered collections of the former (gene sets, in this case).

2. Role model.

2.1. *Model.* Following the description in Newton et al., we have a finite number of *parts* p and a finite number of *wholes* w , where each whole is an unordered set of parts. The incidence matrix $I = (I_{p,w})$ is determined from external knowledge, where $I_{p,w} = 1$ if and only if $p \in w$. The intended correspondence is that genes are parts and gene sets (i.e., functional categories) are wholes. The matrix I encodes a full collection of gene sets. We will have measured data on the parts and aim to make inference on properties of the wholes.

The experimentally-derived gene list may be viewed as a vector of Bernoulli trials $X = (X_p)$, with $X_p = 1$ if and only if part (gene) p is on the list. First proposed in Bauer, Gagneur and Robinson (2010), the role model describes the joint distribution of X in terms of latent binary (0/1) activity variables $Z = (Z_w)$ and by part-level activities induced by them: $A_p = 1$ if $Z_w = 1$ for any w with $p \in w$ or, equivalently,

$$(2.1) \quad A_p = \max_{w: p \in w} Z_w.$$

This conveys the simple assumption that a part is activated if it is in any whole that is activated. For false-positive and true-positive parameters $\alpha, \gamma \in (0, 1)$, with $\alpha < \gamma$, the model for X entails mutually independent components (conditionally on latent activities), with

$$(2.2) \quad X_p \sim \text{Bernoulli} \begin{cases} \alpha, & \text{if } A_p = 0, \\ \gamma, & \text{if } A_p = 1. \end{cases}$$

Simply, activated parts (i.e., those with $A_p = 1$) are delivered to the list at a higher rate than inactivated parts. A key feature of the model is that a part (gene) is activated by virtue of any one of its functional *roles*; this implies that a gene may be activated and yet be part of a functional category that is inactivated, which is in contrast to most other gene-set inference methods [e.g., Goeman and Bühlmann (2007); Barry, Nobel and Wright (2008); Sartor, Leikauf and Medvedovic (2009)] and which provides for a fully multivariate analysis of the gene list. In Bauer, Gagneur and Robinson (2010) it is further assumed, for the sake of Bayesian analysis, that uncertainty in whole-level activities is represented with a single rate parameter $\pi \in (0, 1)$:

$$(2.3) \quad Z_w \sim_{\text{i.i.d.}} \text{Bernoulli}(\pi).$$

Taken together, the model (2.2) and the prior (2.3) determine a joint posterior for Z given X . The R package MGSA (model-based gene-set analysis) reports MCMC-computed marginal posterior probabilities $P(Z_w = 1|X)$, also integrating uncertainty in the system parameters α, γ and π , and thus provides a useful ranking of the wholes [Bauer, Robinson and Gagneur (2011); R Development Core Team (2011)].

In addition to the system incidence matrix I , a useful data structure for computations turns out to be the bipartite graph \mathcal{G} , having whole nodes and part nodes, and an edge between w and p if and only if $I_{p,w} = 1$ (i.e., iff $p \in w$).

2.2. Activation hypothesis. As defined above, the role model allows that a whole can be inactive while all of its parts are active. This can happen because of overlap among the wholes. Specifically, if w is contained in the union of other wholes $\{w'\}$, then all $Z_{w'} = 1$ will force $A_p = 1$ for all $p \in w$, regardless of the value of Z_w . This rather odd situation calls into question the meaning of *active* and what we might realistically expect can be inferred from data. Indeed, the issue is related to identifiability of the activity vector Z , since it shows that distinct Z vectors may produce the same part-level activity vector $A = (A_p)$. (In the case above, switching Z_w from 0 to 1 does not change A .) The mapping $Z \rightarrow A$ given by (2.1) is not necessarily invertible, depending on the system as defined in I . Lack of identifiability would not necessarily create difficulty in a Bayesian analysis, however, in the present case we are specifically interested in inferring the activity states of the gene sets and prioritizing these sets, and so it stands to reason that we ought to confer a real, if still only model-based, meaning on the activities.

When *activity* is defined more fully, there is a simple solution to the problem. The *activation hypothesis* asserts that a set of parts is active if and only if all parts in the set are active. It was shown previously (Newton et al.) as follows:

THEOREM 2.1. *Under the activation hypothesis (AH), the mapping $Z \rightarrow A$ defined by*

$$A_p = \max_{w: p \in w} Z_w$$

is invertible, with inverse $A \rightarrow Z$

$$Z_w = \min_{p: p \in w} A_p.$$

The inverse mapping is simply that a whole is inactive if and only if any of its parts is inactive. So the odd case at the beginning of the section cannot occur under AH; if all parts are active, then $Z_w = 1$ must hold. Further, with parameters α and γ fixed, the Z vector is identifiable under AH, since different Z vectors necessarily give different probability distributions to data X .

The first contribution of the present work is to show that the activation hypothesis is equivalent to a set of linear inequality constraints on the activity variables. The finding is useful for posterior inference computations. We prove in Section 7 the following:

THEOREM 2.2. *AH holds if and only if all of the following hold:*

1. $Z_w \leq A_p$ for all p, w with $p \in w$;
2. $A_p \leq \sum_{w: p \in w} Z_w$ for all p ;
3. $\sum_{p: p \in w} (Z_w - 2A_p + 2) \geq 1$ for all w .

Evidently, the i.i.d. Bernoulli prior (2.3) does not respect AH in the sense that vectors Z which violate AH have positive prior probability. In simple systems such violation may be innocuous. We provide evidence that in the complex systems such as GO, this violation creates a substantial loss of statistical efficiency. We note first that alternative prior specifications are available that respect AH. A simple one is to condition prior (2.3) on the AH event, namely,

$$(2.4) \quad P(Z = z) = \left(\frac{1}{c}\right) \pi^{\sum_w z_w} (1 - \pi)^{\sum_w (1 - z_w)} \quad \text{if } z \text{ satisfies AH,}$$

otherwise $P(Z = z) = 0$, where c is the probability, in prior (2.3), that Z satisfies AH, and z is a vector of binaries representing a possible realization of Z . In other words, with subscript “1” for the i.i.d. prior (2.3) and “2” for prior (2.4), we have $P_2(Z = z) = P_1(Z = z | \text{AH})$. Upon conditioning, the (Z_w) are not necessarily either mutually independent or identically distributed.

3. Statistical properties. The role of the prior distribution in Bayesian analysis has surely been the subject of considerable debate. On the one hand, it helps by regularizing inference, especially in high dimensions. On the other hand, data need to work against it to produce inferences that trade off empirical characteristics with prior assumptions. A fact of relevance to the present problem is that gene-list data must work against either prior [(2.3) or (2.4)] to deliver an inferred list of activated gene sets. For two Bayesian analysts, one using prior (2.3) and the other using prior (2.4), the true state is ascribed different prior mass. The ratio of these masses, ρ , represents the extra effort needed to be done by the data to overcome prior (2.3) compared to prior (2.4):

$$(3.1) \quad \rho = \frac{P_2(Z = z_{\text{true}})}{P_1(Z = z_{\text{true}})} = \frac{P_1(Z = z_{\text{true}}|\text{AH})}{P_1(Z = z_{\text{true}})} = \frac{1}{P_1(\text{AH})} \geq 1.$$

Here we have used the particular structure of prior (2.4) and also the assumption that z_{true} satisfies AH. If z_{true} did not satisfy AH, the target of inference would be beyond the realm of any gene-level data set to estimate, owing to lack of identifiability. Indeed, it is difficult to see what meaning could be ascribed to z_{true} in that case. The observation to be gained from (3.1) is that the probability of AH under the i.i.d. prior affects the efficiency of inference. In systems where that probability is very small, there is reason to believe that improved inferences are possible. As to the precise effect of ignoring AH, that depends on the particular system I , the true activation state, and the system parameters α and γ . What our initial investigation finds is that a truly activated whole w may tend to have $P_1(Z_w = 1|X)$ smaller than $P_2(Z_w = 1|X)$, and if so the P_1 inference is too conservative.

Whether or not AH holds for a given state Z may be assessed by calculating the part-level activities A and then checking Theorem 2.2. Alternatively, we consider whole-level *violation* variables (V_w). These Bernoulli trials are defined as follows:

$$(3.2) \quad V_w = \begin{cases} 1, & \text{if } Z_w = 0 \text{ and if for all } p \in w \text{ there exists } w' \\ & \text{with } p \in w' \text{ and } Z_{w'} = 1, \\ 0, & \text{otherwise.} \end{cases}$$

The probability, under P_1 , that Z satisfies AH is equivalent to the probability of no violations, that is,

$$(3.3) \quad P_1(\text{AH}) = P_1(V_w = 0, \forall w),$$

and so AH probability might be approachable by considering the violation variables. Except in stylized examples, we do not expect these variables to be mutually independent; indeed, they may have a complicated dependence induced by overlaps of the wholes and, hence, direct calculation of $P_1(\text{AH})$ is intractable. However, the expectations of V_w are readily computable for a given system, either by Monte Carlo or by a more sophisticated algorithm [Wang et al. (2014)]. Considering the Chen–Stein result for Poisson approximations, we conjecture that $-\log P_1(\text{AH})$ is approximately equal to $E_1(\sum_w V_w)$, though we have not been

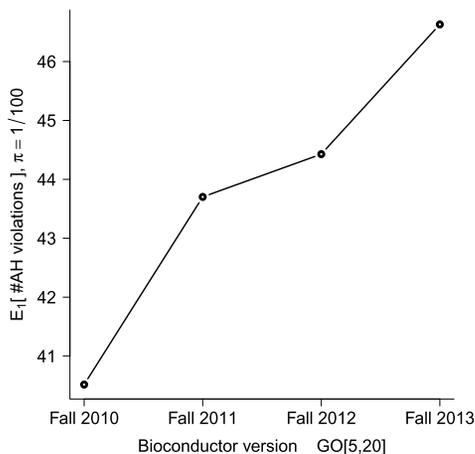


FIG. 1. Expected number of sets that violate the activation hypothesis (AH) for four recent versions of Gene Ontology (GO), considering sets holding between 5 and 20 genes, taken on the *i.i.d.* Bernoulli prior. Calculations are done at $\pi = 1/100$. Respectively, these systems contain 3591, 4096, 4449 and 4772 gene sets, and correspond to versions of *org.Hs.eg.db*.

able to guarantee an error on this approximation [cf. Arratia, Goldstein and Gordon (1990)].

Figure 1 charts the expected value $E_1(\sum_w V_w)$ over four recent versions of Gene Ontology, for $\pi = 1/100$. For concreteness it focuses on GO terms holding between 5 and 20 genes (for which an exact calculation of the expectation is feasible), though the key finding is not sensitive to that restriction, as evidenced by Monte Carlo computations (not shown). As one might expect by the increasing density and complexity of GO, the expected number of AH violations increases. This may very well reflect the fact that $P_1(\text{AH})$ is decreasing over time, which indicates to us that ignoring AH is becoming an ever greater problem for gene-set analysis.

In terms of modeling assumptions, there is no additional cost to accounting for AH in the Bayesian analysis; the cost is purely computational, since inference must now deal with the constraints imposed by AH on the space of latent activities. The next sections describe two computational advances that address the problem.

4. Decoding functional signals via constrained optimization.

4.1. *MAP via ILP.* Decoding a discrete signal is frequently accomplished by algorithms that compute the parameter state having the highest posterior mass: the maximum a posteriori (MAP) estimate. Although limited as a posterior summary, the MAP estimate may contain useful multivariate information [e.g., Carvalho and Lawrence (2008)]. Our representation of model-based gene-set analysis reveals

that under model (2.2) and prior (2.4), the log posterior is linear in the joint collection of whole and part activity variables. This log posterior is

$$(4.1) \quad \begin{aligned} l(Z, A) = & \sum_w \{Z_w \log(\pi) + (1 - Z_w) \log(1 - \pi)\} \\ & + \sum_p \{A_p [x_p \log(\gamma) + (1 - x_p) \log(1 - \gamma)] \\ & + (1 - A_p) [x_p \log(\alpha) + (1 - x_p) \log(1 - \alpha)]\}, \end{aligned}$$

where x_p is the realized value of the gene-list indicator X_p , and α , γ , and π are system parameters, which are considered fixed in the present calculation. Considering Theorem 2.2, finding the MAP estimate (\hat{Z}, \hat{A}) amounts to maximizing a linear function in discrete variables subject to linear inequality constraints. As such, it fits naturally into the domain of *integer linear programming* (ILP), an active subfield of optimization. Our computations take advantage of ILP software available in the GNU Linear Programming Kit through its interface with R.⁶ We employed a series of basic code checks to assure our implementation worked in: (1) simple examples where the MAP estimate is computable by other means; and (2) limiting situations where X_p was Binomial having a high sample size, and thus where the MAP estimate must converge to the true activity state.

The reconstruction \hat{Z} obtained through this optimization holds an estimate of the activated and inactivated gene sets. We refer to the overall method as the *multi-functional analyzer* (MFA), and specifically MFA-ILP to refer to the posterior mode computed by ILP. We note that by invertibility of the mapping $Z \rightarrow A$ under AH, the log-posterior l could be expressed either as a function of Z only or as a function of A only, however, in neither reduced case would l be linear in the input variables. Moreover, in neither reduced case could the constraints be expressed as linear inequality constraints. By expanding the domain we formulate the constrained optimization as an integer linear program.

4.2. *Numerical experiments.* In each experiment reported below we represented a system with a parts-by-wholes incidence matrix I ; we fixed the false-positive rate $\alpha = 1/10$ and the true positive rate $\gamma = 9/10$. We simulated 100 gene-lists X from model (2.2), each time using a simulated activity vector Z . For methods comparison, we applied the following: (1) the commonly used Fisher exact test for enrichment of each gene set in the data X [Khatri, Sirota and Butte (2012)], (2) MGSA (version 1.7.0), and (3) MFA-ILP. We allowed both model-based methods to know the system parameter settings. To evaluate performance, we calculated specificity, sensitivity, and precision of the estimated activity vector \hat{Z} for the true activity vector Z by averaging over the 100 replicates.

⁶See www.gnu.org/software/glpk and cran.R-project.org/package=Rglpk.

TABLE 1

Simulation comparison of three gene-set methods, a case of low overlap among gene sets: Tabulated are mean values from 100 simulated data sets. On average 7.3 truly activated sets occur

Method	Predicted # active	Sensitivity	Specificity	Precision
MFA-ILP	7.4	0.963	0.997	0.958
Fisher (cut-off = 0.05)	5.9	0.790	0.998	0.966
Fisher (cut-off = 0.1)	6.8	0.873	0.996	0.948
MGSA (cut-off = 0.5)	7.2	0.954	0.998	0.968

Experiment 1: Low overlap. Initially, I had size 300 genes (parts) by 100 gene sets (wholes). We randomly picked 5 and 10 parts for each whole in columns 1–50 and 51–100, respectively. Then we removed parts not contained by any whole, leaving a 296×100 incidence matrix. We sampled Z from prior (2.3) and then projected it onto AH by constructing $A_p = \max_{w: p \in w} Z_w$ and then updating $Z_w = \min_{p: p \in w} A_p$. All methods exhibit similar operating characteristics in this case (Table 1).

Experiment 2: Higher overlap and parent-child structure. Initially, I had size 300 parts by 105 wholes. From column 1 to column 20, each column has 20 parts, of which 15 parts are in common with each other and 5 parts are randomly selected from the other parts; column 21 has 10 parts which are randomly picked from the 15 common parts shared by columns 1–20. Thus, columns 1–20 have a lot of overlaps and column 21 is a child of columns 1–20. Similarly, we built columns 22–42, 43–63, 64–84 and 85–105. The common 15 parts in each column combination are all different. Then parts not contained by any whole were removed, which resulted in a 265×105 incidence matrix. We activated wholes by sampling one whole from columns 1–20, 22–41, 43–62, 64–83 and 85–104 as activated, respectively, and projected onto AH as above.

Table 2 exhibits properties of three methods in relatively complicated system just defined. The univariate Fisher test tends to select the wholes with a high corre-

TABLE 2

Simulation comparison of three gene-set methods, a case of higher overlap among gene sets: Tabulated are mean values from 100 simulated data sets. On average there are 10.1 truly activated sets in this case

Method	Predicted # active	Sensitivity	Specificity	Precision
MFA-ILP	10.2	0.975	0.997	0.993
Fisher (cut-off = 0.05)	104.2	0.996	0.008	0.096
Fisher (cut-off = 0.1)	104.8	0.996	0.002	0.096
MGSA (cut-off = 0.5)	5.5	0.490	0.995	0.920

lation (overlap) with the truly activated wholes, which results in high sensitivity but low specificity (or precision). The extra activation calls correspond to spurious associations that the multivariate, model-based approaches are able to recognize. The MGSA method often fails to discover truly activated wholes, which corresponds to a reduced sensitivity. The small *child* wholes tend to be missed by MGSA in this case. The proposed MFA-ILP method is right on target.

4.3. *ILP for large systems.* Large systems strain unaided ILP computation, but the special structure of the gene-set problem allows for several refinements.

Shrinking I. Up to a constant, the objective function in (4.1) may be expressed

$$l(Z, A) = c_1 \sum_w Z_w + c_2 \sum_{p \in P^-} A_p + c_3 \sum_{p \in P^+} A_p,$$

where

$$c_1 = \log \pi - \log(1 - \pi),$$

$$c_2 = \log(1 - \gamma) - \log(1 - \alpha),$$

$$c_3 = \log \gamma - \log \alpha$$

and where P^- and P^+ denote the observed inactivated and activated parts, respectively. That is, $p \in P^-$ if $x_p = 0$ and $p \in P^+$ if $x_p = 1$. By assumption (2.2), $\alpha < \gamma$ and so $c_2 < 0$ and $c_3 > 0$. If we further insist that $\pi < 1/2$, then $c_1 < 0$ also. In some cases we can know which \hat{Z}_w and \hat{A}_p must equal 0 in the optimal solution, and if so we can remove these variables from the system prior to implementing ILP. For each whole w denote $P_w^+ = w \cap P^+$ and similarly $P_w^- = w \cap P^-$, and define $W^* = \{w : c_1 + c_3 \sum_{p \in P_w^+} 1 < 0\}$. Clearly, those wholes containing no reported parts are in W^* , but there may be others. We prove in Section 7 that if W^* is not empty, then we may be able to shrink the system prior to solving the constrained optimization problem via ILP.

THEOREM 4.1. *Suppose $\pi < 1/2$ and let w_0 denote an element of W^* . If there exists $p_0 \in w_0$ such that $\{w : p_0 \in w\} \subset W^*$, then $\hat{Z}_{w_0} = \hat{A}_{p_0} = 0$.*

Letting W_0 and P_0 denote wholes and parts for which the optimal solution is known (in advance of computation), we may remove these from the incidence matrix I , effectively shrinking it. The amount of shrinkage may be dramatic, but it depends on the observed data x , the system I , and system parameters α , γ and π . When α is small and γ is large, the effects may be minimal.

A sequential approach. In the unlikely event that the system matrix I is separable into blocks of wholes that do not overlap between blocks, then ILP may be

applied separately to these distinct blocks in order to identify the MAP activities. We do not expect this separability in GO or related systems, but we can take advantage of size variation of the wholes and work sequentially from small ones to larger ones. As an example, let S_{10} denote the sets containing no more than $n.up = 10$ genes. In order to obtain the optimal solution for the full problem, we start from the sub-matrix $I.10$ obtained by extracting these sets from I . Suppose Z_{10}^* is the MAP solution based on the data for $I.10$, and use notation S_{10}^* to denote the active sets in S_{10} as inferred by Z_{10}^* . We aim to find the optimal solution Z_{11}^* for $I.11$ using what has already been computed in the smaller system. Denote the newly added sets in $I.11$ by S_{10}^{11} (i.e., the sets containing exactly 11 genes). We just need to consider the sets with the possibility being active in the optimal solution on $I.11$. First of all, S_{10}^* and S_{10}^{11} should be included, in the case we have no any other prior knowledge about Z_{11}^* . Second, by the 3rd AH inequality (Theorem 2.2), any set in $S_{10} \setminus S_{10}^*$ which is a subset of some set in S_{10}^{11} , denoted by D , should also be included. But these sets already considered are not enough. Actually, for each set w_1 in S_{10}^{11} , we need to check whether there exists some set w_2 in $S_{10} \setminus (S_{10}^* \cup D)$ satisfying

$$(4.2) \quad c_1 + c_2 \sum_{p \in P^- \cap P_{w_1}^{w_2}} A_p + c_3 \sum_{p \in P^+ \cap P_{w_1}^{w_2}} A_p > 0,$$

where $P_{w_1}^{w_2}$ denote the set of genes contained by w_2 and not by w_1 . We do this since each set in S_{10}^{11} may be active in the optimal solution Z_{11}^* , and we need to check whether some sets in S_{10} should be activated toward maximizing the objective function. We denote the sets in $S_{10} \setminus (S_{10}^* \cup D)$ satisfying the condition (4.2) by E . Finally, by the 3rd AH inequality (Theorem 2.2), any set in $S_{10} \setminus (S_{10}^* \cup D \cup E)$ which is a subset of some set in E , denoted by F , should also be included. Thus, we need to run the ILP on the incidence matrix only for $S_{10}^* \cup S_{10}^{11} \cup D \cup E \cup F$, instead of $I.11$. Hence, we obtain a sequential approach to solve the full ILP problem from a sequence of smaller problems. Examples show this is feasible in GO for subsystems holding sets of up to 50 genes, without excessive computational burden.

5. Posterior sampling.

5.1. *Penalized MCMC.* To obtain a sample from the posterior distribution defined by prior (2.4) and model (2.2) in which the whole activity variables $Z = (Z_w)$ have positive probability only when Z satisfies AH, we design a Markov chain to run within the unconstrained space according to a penalized posterior:

$$(5.1) \quad \tilde{l}(Z) = l(Z, A) - \lambda \sum_w V_w,$$

where $l(Z, A)$ is defined in (4.1), V_w is the violation indicator (3.2), and $\lambda \geq 0$ is a tuning parameter. The desired sample is obtained by discarding any sampled states

that do not satisfy AH. Note that there are no violations ($\sum_w V_w = 0$) for Z that satisfy AH, so that $\tilde{l}(Z) = l(Z, A)$ in this case and the conditional log posterior distribution under $\tilde{l}(Z)$ restricted to AH is identical to the target log posterior distribution. Increasing the tuning parameter λ increases the probability of AH in the larger state space, which is essential for efficient sampling when this probability is small. We find that penalizing the log posterior within the unconstrained space leads to a conditional sampler that mixes well in the constrained space, where our previous attempts to constrain move types were less successful.

It is helpful to visualize the Markov chain as operating by changing colors on the node-colored bipartite graph \mathcal{G} having whole nodes and part nodes, with an edge between a whole node w and part node p if and only if $p \in w$, and where the coloring of the whole nodes $\{w\}$ and part nodes $\{p\}$ match the activities Z and A , respectively. It is useful in assessing the state of the Markov chain to associate with each node a count $n(\cdot)$ of its active connected neighbors in \mathcal{G} . $A_p = 1$ if and only if $n(p) > 0$ and $V_w = 1$ if and only if $Z_w = 0$ and $n(w) = \text{deg}(w)$, the number of part nodes $p \in w$.

The Markov chain proceeds by selecting at random a whole node w and proposing a color swap (a change in the status of the activity variable, $Z_w^* = 1 - Z_w$) for this node.⁷ This proposed change can, but need not, affect the activities of parts contained in this whole. When $Z_w^* = 1$, the active neighbor counts $n(p)$ increase by 1 for each $p \in w$. If A_p changes from 0 to 1, then each node w' that contains p (including w) gains an additional active neighbor and $n(w')$ increases by 1. This increase could cause a violation if p were the only remaining inactive neighbor of an inactive w' , causing $V_{w'}$ to change from 0 to 1. If node w were in violation before this proposal, activating it would eliminate the violation. Similarly, when $Z_w^* = 0$, the active neighbor counts $n(p)$ decrease by 1 for each $p \in w$. If this decrease is from 1 to 0, then the activity A_p changes from 1 to 0 as well and all of the whole nodes w' connected to p would lose an active neighbor, $n(w')$ decreasing by 1. If the whole node w' had been in violation, this change would eliminate the violation with $V_{w'}$ changing from 1 to 0.

Careful accounting of the changes to a few key counts allows for quick calculation of the change in $\tilde{l}(Z^*)$ and subsequent acceptance or rejection of the proposal by Metropolis–Hastings. The log posterior $\tilde{l}(Z^*)$ is a function of α , γ , π , the penalty λ and the counts of the numbers of active and inactive whole nodes [$\sum_w Z_w$ and $\sum_w (1 - Z_w)$, resp.], the number of whole nodes in violation ($\sum V_w$), the numbers of active part nodes with realized values 1 and 0 [$\sum_p A_p x_p$ and $\sum_p A_p (1 - x_p)$, resp.], and the numbers of inactive part nodes with realized values 1 and 0 [$\sum_p (1 - A_p) x_p$ and $\sum_p (1 - A_p) (1 - x_p)$, resp.].

⁷Efficiency gains may be possible by using a nonuniform sampler, for instance, depending on the set size or the number of reported parts in the whole, though we use a uniform proposal in the present work.

5.2. Numerical experiment. To assess the performance of MFA-MCMC, we simulated gene-list data according to the role model in the *D. melanogaster* genome, following the scheme presented in Bauer, Gagneur and Robinson (2010). Briefly, we used 3275 GO terms annotating between 5 and 50 fly genes, according to version 2.14.0 of Bioconductor package `org.Dm.eg.db`. In each simulation run, a number of GO terms were activated and then a gene list was constructed from independent Bernoulli trials depending on the activation states and settings of false-positive and false-negative error rates. Figure 2 shows receiver-operating (ROC) curves and precision-recall curves for two parameter settings, based on 100 simulated gene lists in each setting. Selection to the reported set list is based on thresholding the marginal posterior probability (MGSA, MFA-MCMC) or the

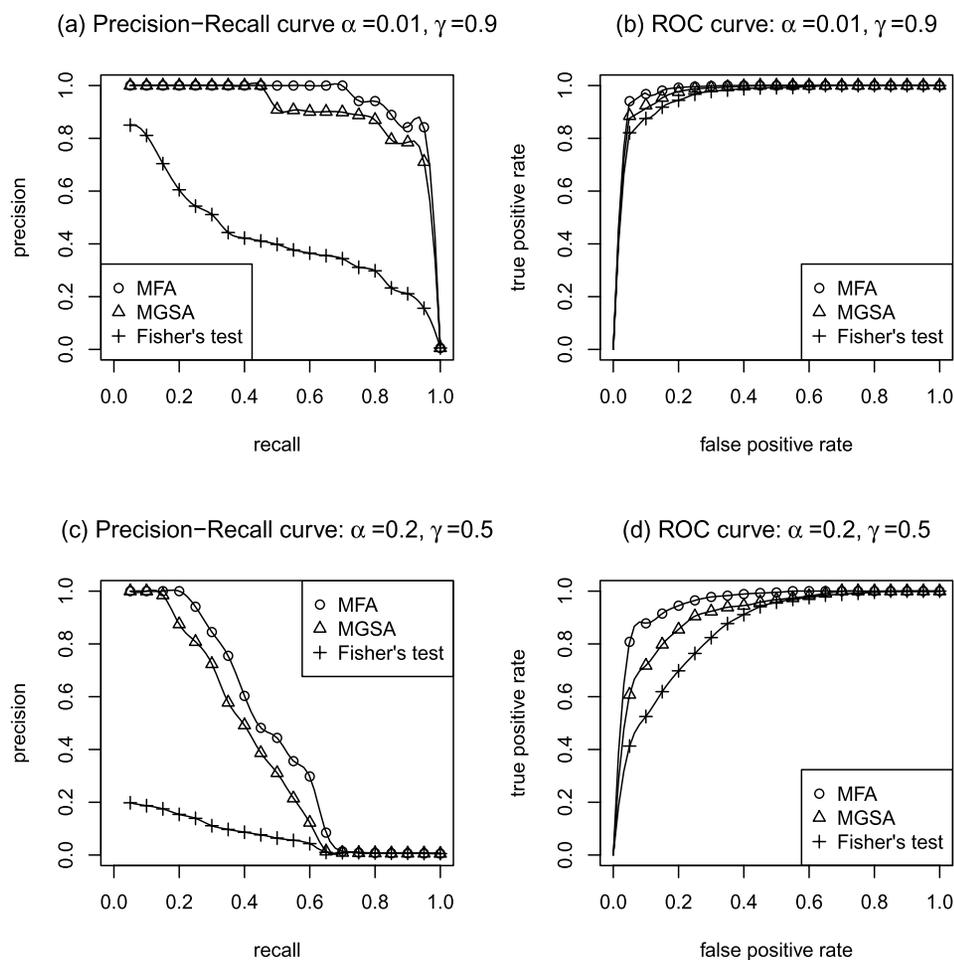


FIG. 2. Operating characteristics of MFA-MCMC, MGSA, and Fisher's test based on simulating the role model in the *D. melanogaster* genome.

p -value (Fisher). Evidently, MGSA and MFA-MCMC are accurate and show similar behavior when error rates are low, though MFA-MCMC shows improved precision and sensitivity in more difficult settings.

In subsequent calculations we deploy both MAP estimation (MFA-ILP) and MCMC sampling on each data set in order to infer wholes that are probably activated. For MCMC, we use 10^7 sweeps, burn-in of 10^6 , and $\lambda = 5$, which causes about one third of the states to satisfy AH. MFA-ILP gives a summary functional decoding of the gene list. Posterior probabilities from the MCMC computation provide a measure of confidence in the inferred sets and also highlight notable non-MAP sets. Fisher's test is the default univariate method for gene-list data; we include it for comparison, even though the hypotheses it tests are different from the activation states assessed by MFA and MGSA.

6. Examples.

6.1. *Genes implicated in type 2 diabetes (T2D)*. From a large-scale genome-wide association study (GWAS) involving more than 34,000 cases and 114,000 control subjects, 77 human genes have been implicated as affecting T2D disease susceptibility [Morris et al. (2012), Supplementary Table 15, primary list]. To assess the functional content of this gene list, we applied MFA, MGSA, and simple enrichment via Fisher's exact test, all in the context of 6037 gene ontology terms, each annotating between 5 and 50 genes.⁸ Here and in other examples we took advantage of available information on likely false positive (α) and false negative ($1 - \gamma$) error rates at the gene level. Using the fitted mixture model from Morris et al., we estimated $\alpha = 0.00019$ and $\gamma = 0.02279$ for this large-scale GWAS [details in Wang et al. (2014)].

Figure 3 summarizes the application of MFA, MGSA, and Fisher's test to this example. Table 3 reports those gene sets inferred by MFA-ILP to be activated in T2D. Tables S1–S4 provide further information for comparison of MFA with MGSA and Fisher's test [Wang et al. (2014)]. The example illustrates features we see repeatedly with these methods. Sets identified by Fisher's test tend to overlap substantially, reflecting the univariate nature of the approach; both MGSA and MFA-ILP alleviate this redundancy problem, but MGSA finds fewer sets than MFA-ILP. As expected, each of the 11 sets inferred by the ILP algorithm (i.e., the MAP estimate of the activated sets) has high marginal posterior probability of activation (P.MFA). Furthermore, MFA-ILP is able to explain more of the gene level findings than the other methods, as indicated by the number of genes that are both in the reported gene list (T2D) and are in at least one of the gene sets inferred to be activated (*coverage*). It does this without increasing the mis-coverage, which is

⁸These 6037 terms annotate a total of 10,626 genes; among the 77 T2D-associated genes, 58 are in this *moderately annotated* class.

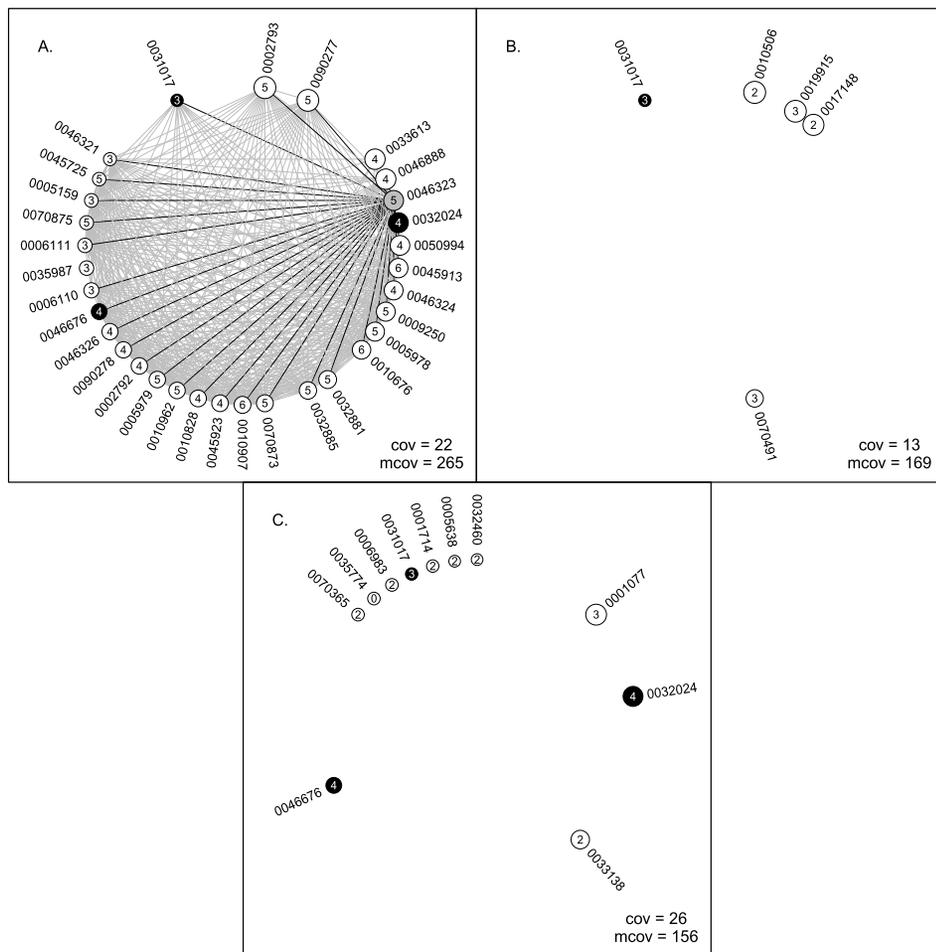


FIG. 3. GO terms inferred by three methods (A, Fisher's test; B, MGSA; C, MFA-ILP) to be activated to explain the type 2 diabetes associated genes (58 of the 77 are annotated to the 6037 GO terms used). Panels are linked by the location of sets; three sets (black) were identified by multiple methods. GO ID is on the outer rim, and sets are ordered by size as in Tables 3 and S1–S3, from largest (noon) and decreasing clockwise; numbers inside circles indicate the number of T2D genes annotated to that set; a line connects two sets if their intersection contains at least one T2D-associated gene. The total number of T2D genes explained by the identified sets (coverage), lower right, as a subset of the 58 starting genes. Similarly, the mis-coverage (mco) counts the number of genes within the inferred active sets that are not on the observed gene list. GO:0046323 (glucose import) and its relations are highlighted in panel A (grey circle, black lines).

the number of non-T2D genes within the inferred active sets. Figure 3 summarizes the sets found by these methods and reports these coverages.

An interesting set in this case is the GO term *glucose import* (GO:0046323), for which the proportion 5/41 of observed T2D genes is very high (small Fisher *p*-value), but there is very small posterior activation probability according to MFA.

TABLE 3

MFA results in type 2 diabetes (T2D) example: 11 GO terms are inferred to be active using the ILP algorithm to compute the MAP estimate (rows). Basic statistics on these terms are provided (# T2D-associated genes/set size). The next two columns give the MCMC-computed marginal posterior activation probabilities for these terms, both using MGSA and MFA, the constrained alternative. The final column holds the Benjamini–Hochberg adjusted Fisher-test p -value. All calculations start with 6037 GO terms (those annotating between 5 and 50 human genes) that together annotate 10,626 human genes. Of the 77 total T2D genes, 58 have at least one annotation to these 6037 GO terms. The inferred gene sets cover 26 of these 58 T2D genes

Gene set (GO term)	Statistics	P.MFA	P.MGSA	Fisher
RNA polymerase II core promoter...	3/45	0.517	0.028	0.161
Positive regulation of insulin secretion	4/41	0.964	0.372	0.016
Positive regulation of peptidyl-serine...	2/35	0.537	0.096	0.756
Negative regulation of insulin secretion	4/23	0.996	0.201	0.003
ER overload response	2/9	0.398	0.159	0.102
Positive regulation of insulin secretion...	0/9	0.964	0.002	1.000
Hepatocyte differentiation	2/9	0.316	0.016	0.102
Endodermal cell fate specification	2/8	0.596	0.036	0.091
Exocrine pancreas development	3/8	0.946	0.600	0.003
Negative regulation of protein...	2/5	0.420	0.101	0.051
Lamin filament	2/5	0.790	0.400	0.051

That is because the 5 genes are explained more easily as parts of three other terms in the MAP estimate that have yet other genes supporting their activation.

A second curious case is GO:0035774, a small term (9 genes) to do with regulation of insulin secretion. None of these genes was reported to be involved in T2D, however, the set is fully contained in a parent set which is inferred to be activated by MFA-ILP. As the calculation respects AH, all subsets of activated sets are activated. This may be a set-level false-positive call, as none of the contained genes was reported to be T2D associated. MFA favors the explanation that each of the 9 noncalls was a gene-level false negative, finding the weight of evidence supporting that interpretation. When we recall that the gene-level false-negative rate is almost 98% (following the mixture calculation from Morris et al.), this assessment seems plausible. We note that for the sake of further simplification of output, it is reasonable to suppress any such subsets from primary tabulations [see trimming algorithm, Wang et al. (2014)].

6.2. *RNA interference and influenza-virus replication.* In a meta-analysis of four genome-wide studies of influenza virus, Hao et al. (2013) reported that 984 human genes had been detected by RNA interference as possibly being associated with viral replication. As in the T2D example, we compared MFA with MGSA and Fisher's test on this gene list using 6037 GO terms annotating between 5 and 50 human genes. Among the 984 influenza-involved genes, 683 are annotated to

process, which annotates 48 human genes, 12 of which were in the 984-list of influenza involved genes (highlighted in Figure 4). The Benjamini–Hochberg corrected Fisher p -value is 0.017, and so the term would be considered enriched in the standard analysis; it is also inferred to be active by MGSA (posterior probability 0.897). But it is not in the MAP estimate by MFA-ILP and its posterior activation probability is 0.000. Now 10 of these 12 influenza-involved GO:0032434 genes are part of the *child* term GO:0032436, *positive regulation of proteasomal ubiquitin-dependent protein catabolic process*, a set of size 31 genes. Both terms are found by the univariate Fisher procedure, exemplifying the redundancy issue; the more specific term GO:0032436 is identified by MFA-ILP and has high posterior activation probability. The MFA calculation favors the explanation whereby the smaller set is activated; this fails to cover two of the 12 GO:0032434 influenza genes, but it also simplifies the explanation of nonlisted genes in that set. If the larger set (GO:0032434) is activated, then we have a lot more mis-covered genes, that is, those in the set but undetected by RNAi [$15 = (48 - 12) - (31 - 10)$]. With this example, it may be that the more specific *positive regulation* term better characterizes the experimental gene list.

In Hao et al., gene set analysis was used to show that the four separate RNAi studies agreed more substantially than was evident by inspecting overlaps among the four gene lists. It was applied separately to the study-specific gene lists, and then the agreement among these four set lists was measured. For both Fisher's test and MGSA-ILP, the among-study set-level agreement was significant according to a simple permutation calibration. Curiously, the agreement by MGSA was not significant by that measure, owing primarily to the fact that very few sets were inferred to be active in the separate studies. The common set-level signal, in conjunction with other forms of meta-analysis, provided evidence that genome-wide RNAi studies have higher false-negative rates than false-positive rates.

6.3. Other issues. The full effects of prior choice in model-based gene-set analysis require further investigation. As to the practical importance of one choice over another, we do not examine the biological distinctions between the inferences produced by different methods. A close reading of the T2D and RNAi case studies above provides an initial indication of how and why reported set lists can differ, but assessing the biological significance of these differences is beyond the present scope. The procedures have distinct statistical properties and MFA more efficiently captures the functional content of the reported gene list in terms of model fit. We point out that the distinctions present themselves when using the relatively complex GO system. Control calculations show that MGSA and MFA give essentially the same results when applied to the less complex KEGG system [Figures S1 and S2, Wang et al. (2014)].

In our comparisons we used MGSA to obtain an estimate of the hyperparameter π , which affects the overall rate of set activity. In order to control the comparison, we used the same numerical value of π in the MFA calculations (in

both MFA and MGSA we fixed the other parameters α and γ at externally derived values). Further improvements of MFA may be possible using alternative estimates of π . Other model elaborations which could be useful in some applications include extending MFA beyond binary X_p and allowing dependence in gene level measurement errors.

Compute times for MFA depend on the size and content of the gene list, the incidence matrix I , and the model parameters: MFA-ILP used 2.5 CPU hours for T2D, and 23 CPU hours for RNAi;⁹ less time was required for MFA-MCMC (20 and 45 CPU minutes, resp.).

We have argued that temporal changes in GO reflect an increase in the complexity of the functional record that justifies the more refined prior distribution used in MFA (Figure 1). These changes also tell us that the results of a given analysis naturally depend on the GO version, since the sets involved and the annotations of genes to sets continue to evolve. To assess the version effect, we applied MFA to four recent GO versions (2010–2013) in both the T2D and RNAi examples, and using sets annotating between 5 and 50 human genes (GO[5:50]). Table 4 records how similar are time-adjacent versions of GO as well as how similar are results of MFA. The changes in GO reflect new terms and new annotations, as well as sets moving in and out of GO[5:50] as more genes become annotated. Against this

TABLE 4

Version effects: Tabulated are similarity scores comparing set structure and inferred active sets over time-adjacent versions of GO[5:50]. The collections of sets annotating between 5 and 50 human genes fluctuate over recent versions of GO. Respectively, in the four most recent fall versions of Bioconductor, the collections contain 4830, 5546, 6037 and 6488 gene sets. The first row shows the Jaccard index (size of intersection over size of union) comparing subsequent versions of these gene-set collections. In addition to the collections changing, the annotations recording which genes are in which sets also change over time. The second row measures similarity of the sets of annotations. Subsequent rows show similarity of reported lists of gene sets in the two main examples. In this comparison, the set is reported if it is in the MAP estimate by MFA-ILP and if its marginal posterior probability exceeds a threshold (50% or 80%). Inferred active sets depend to some extent on the GO version in view; setting stronger marginal posterior thresholds reduces the false-discovery rate and reduces the version effect

	2010–2011	2011–2012	2012–2013
Sets	0.72	0.82	0.80
Annotations	0.54	0.70	0.66
T2D 50/MAP	0.33	0.60	0.30
T2D 80/MAP	1.00	0.75	0.60
RNAi 50/MAP	0.72	0.70	0.88
RNAi 80/MAP	0.90	0.79	0.85

⁹R was running on a 4 × AMD Opteron(TM) Processor 6174 (48 cores) with 128 GB RAM.

substantial evolution of GO we see that MFA results do also change, but that the changes are less the more stringent is the marginal posterior probability cutoff.

A feature of MGSA and MFA is that activation of the set implies activation of all genes in the set. This strict form of nonnull relationship is quite different from many univariate methods, which would claim a set is nonnull if any of its contained genes is nonnull [e.g., the *self-contained* tests of Goeman and Bühlmann (2007)]. It is precisely this relationship, however, that enables multivariate (i.e., mult-set) analysis, as the role model offers a straightforward approach to deal with the complex overlaps in the functional record. We note that the role model allows a weaker interpretation; for instance, we could continue to assert that a gene is activated if it is contained in any activated set, while allowing that only a fraction of activated genes are nonnull. The difference would be in the tabulation of errors (e.g., $X_p = 0$ might not be a false negative when $A_p = 1$) and in the interpretation of α and γ ; the family of joint distributions would be the same. As GO and other repositories record functions of ever more specific gene combinations, it is reasonable to expect that a combination of genes relevant in the cells on test is within the repository. The strict interpretation of activation is parsimonious and is justified when the repository is sufficiently well endowed with relevant sets. We performed a small simulation study, using the T2D data structure, to assess MFA's ability to recover small activated sets. In each of 100 simulated cases, we fixed the repository (GO[5:50]), we randomized the response vector $X = (X_p)$ by appending to the T2D genes a randomly selected small (5–10 genes) set from GO, and we inferred the activated sets using MFA. In 91 cases the appended set was identified as active by MFA-ILP, demonstrating in a limited way the ability of the methodology to recover signals represented in the repository.

7. Proofs.

7.1. *Proof of Theorem 2.2.* Relative to all the sets and parts in the system I , we say AH holds if and only if $A_p = \max_{w: p \in w} Z_w$ for all p and $Z_w = \min_{p: p \in w} A_p$ for all w . Recall that all A_p and Z_w are binary, in $\{0, 1\}$. The first condition $A_p = \max_{w: p \in w} Z_w$ implies $A_p \geq Z_w$ for all w with $p \in w$; that A_p achieves the max of the Z_w 's has to account for the possibility that $A_p = 1$ when all $Z_w = 0$, but this is covered by having $A_p \leq \sum_{w: p \in w} Z_w$. Thus, the condition $A_p = \max_{w: p \in w} Z_w$ is equivalent to the first two constraints in Theorem 2.2.

To address the second condition, that $Z_w = \min_{p: p \in w} A_p$ for all w , define a new variable

$$T_w = 1 + \sum_{p: p \in w} (A_p - 1),$$

and notice that $T_w = 1$ if and only if $A_p = 1$ for all $p \in w$, otherwise $T_w \leq 0$. Observe that the second condition is equivalent to

$$(7.1) \quad \sum_{p: p \in w} (Z_w - A_p + 1) - T_w \geq 0$$

since if all $A_p = 1$, for $p \in w$, then $T_w = 1$, and Z_w must equal 1 to satisfy (7.1); otherwise, if at least one of the A_p 's equals 0, then $T_w \leq 0$, and noting that the summation in (7.1) is positive confirms the claim. Next, replacing T_w in (7.1) with its definition, we obtain the third stated inequality

$$\sum_{p: p \in w} (Z_w - 2A_p + 2) \geq 1.$$

7.2. Proof of the Theorem 4.1. Compared to $Z_{w_0} = 0$, the possible maximal value added to the objective function by letting Z_{w_0} be 1 is $c_1 + c_3 \sum_{p \in P_{w_0}^+} 1$ (considering parts in $P_{w_0}^-$ may already be activated or w_0 has no parts in P^- , the best case), however, which is negative since $w_0 \in W^*$. So $Z_{w_0} = 0$ is preferred toward maximizing the objective function. Next we need to prove that letting $Z_{w_0} = 0$ and $A_{p_0} = 0$ will not violate the inequalities in Theorem 2.2.

Denote by W_0 and P_0 the sets of w_0 and p_0 satisfying the state in the theorem, respectively. We claim that for each $p_0 \in P_0$, $\{w : I_{p_0, w} = 1\} \subset W_0$. If not, then there exists $w^* \in W^* \setminus W_0$ such that $I_{p_0, w^*} = 1$, so w^* will be in W_0 . This is a contradiction. Thus, $A_{p_0} = \max_{\{w : I_{p_0, w} = 1\}} Z_w = 0$, so the first two AH inequalities are satisfied. It is readily verified that the third inequality is also satisfied.

Acknowledgments. We thank Christina Kendzioriski and anonymous reviewers for comments that helped to guide this research. Both the supplementary material document and a software implementation of MFA are available at <http://www.stat.wisc.edu/~newton/>. An initial version of this manuscript was released as Technical Report #1174, Department of Statistics, University of Wisconsin, Madison. The current version is dated August, 2014.

SUPPLEMENTARY MATERIAL

Supplement: More on role modeling (DOI: 10.1214/14-AOAS777SUPP;.pdf). We provide further details on violation probabilities, on estimating false-positive and true-positive error rates, on preparing data for the ILP algorithm, and on further data analysis findings in the T2D and RNAi examples.

REFERENCES

- ARRATIA, R., GOLDSTEIN, L. and GORDON, L. (1990). Poisson approximation and the Chen–Stein method. *Statist. Sci.* **5** 403–434. [MR1092983](#)
- BARRY, W. T., NOBEL, A. B. and WRIGHT, F. A. (2008). A statistical framework for testing functional categories in microarray data. *Ann. Appl. Stat.* **2** 286–315. [MR2415604](#)
- BAUER, S., GAGNEUR, J. and ROBINSON, P. N. (2010). GOing Bayesian: Model-based gene set analysis of genome-scale data. *Nucleic Acids Res.* **38** 3523–3532.
- BAUER, S., ROBINSON, P. N. and GAGNEUR, J. (2011). Model-based gene set analysis for Bioconductor. *Bioinformatics* **27** 1882–1883.
- CARVALHO, L. E. and LAWRENCE, C. E. (2008). Centroid estimators for inference in high-dimensional discrete spaces. *Proc. Natl. Acad. Sci. USA* **105** 3209–3214.

- GENTLEMAN, R., CAREY, V. J., BATES, D. M., BOLSTAD, B., DETTLING, M., DUDOIT, S., ELLIS, B., GAUTIER, L., GE, Y. et al. (2004). Bioconductor: Open software development for computational biology and bioinformatics. *Genome Biol.* **5** R80.
- GOEMAN, J. J. and BÜHLMANN, P. (2007). Analyzing gene expression data in terms of gene sets: Methodological issues. *Bioinformatics* **23** 980–987.
- HAO, L., HE, Q., WANG, Z., CRAVEN, M., NEWTON, M. A. and AHLQUIST, P. (2013). Limited agreement of independent RNAi screens for virus-required host genes owes more to false-negative than false-positive factors. *PLoS Comput. Biol.* **9** e1003235, 20. [MR3131664](https://doi.org/10.1371/journal.pcbi.1003235)
- KANEHISA, M. and GOTO, S. (2000). KEGG: Kyoto encyclopedia of genes and genomes. *Nucleic Acids Res.* **28** 27–30.
- KHATRI, P., SIROTA, M. and BUTTE, A. J. (2012). Ten years of pathway analysis: Current approaches and outstanding challenges. *PLoS Comput. Biol.* **8** e1002375.
- MATTHEWS, L., GOPINATH, G., GILLESPIE, M., CAUDY, M., CROFT, D., DE BONO, B., GARAPATI, P., HEMISH, J., HERMIAKOB, H., JASSAL, B., KANAPIN, A., LEWIS, S., MAHAJAN, S., MAY, B., SCHMIDT, E., VASTRIK, I., WU, G., BIRNEY, E., STEIN, L. and D'EUSTACHIO, P. (2009). Reactome knowledgebase of biological pathways and processes. *Nucleic Acids Res.* **37** D619–D622.
- MORRIS, A. P. et al. (2012). Large-scale association analysis provides insights into the genetic architecture and pathophysiology of type 2 diabetes. *Nat. Genet.* **44** 981–990.
- NEWTON, M. A., HE, Q. and KENDZIORSKI, C. (2012). A model-based analysis to infer the functional content of a gene list. *Stat. Appl. Genet. Mol. Biol.* **11** Art. 9, 27. [MR2935751](https://doi.org/10.1186/1475-2875-11-27)
- R DEVELOPMENT CORE TEAM (2011). *R: A Language and Environment for Statistical Computing*. R Foundation for Statistical Computing, Vienna, Austria. Available at <http://www.R-project.org/>.
- SARTOR, M. A., LEIKAUF, G. D. and MEDVEDOVIC, M. (2009). LRpath: A logistic regression approach for identifying enriched biological groups in gene expression data. *Bionformatics* **25** 211–217.
- THE GENE ONTOLOGY CONSORTIUM (2000). Gene ontology: Tool for the unification of biology. *Nat. Genet.* **25** 25–29.
- WANG, Z., HE, Q., LARGET, B. and NEWTON, M. A. (2014). Supplement to “A multi-functional analyzer uses parameter constraints to improve the efficiency of model-based gene-set analysis.” DOI:10.1214/14-AOAS777SUPP.

Z. WANG
 Q. HE
 DEPARTMENT OF STATISTICS
 UNIVERSITY OF WISCONSIN, MADISON
 1300 UNIVERSITY AVENUE
 MADISON, WISCONSIN 53706
 USA
 E-MAIL: wangz@stat.wisc.edu
ally00ling@gmail.com

B. LARGET
 DEPARTMENTS OF STATISTICS AND BOTANY
 UNIVERSITY OF WISCONSIN, MADISON
 1300 UNIVERSITY AVENUE
 MADISON, WISCONSIN 53706
 USA
 E-MAIL: larget@stat.wisc.edu

M. A. NEWTON
 DEPARTMENTS OF STATISTICS AND BIostatISTICS
 AND MEDICAL INFORMATICS
 UNIVERSITY OF WISCONSIN, MADISON
 1300 UNIVERSITY AVENUE
 MADISON, WISCONSIN 53706
 USA
 E-MAIL: newton@stat.wisc.edu
 URL: <http://www.stat.wisc.edu/~newton/>