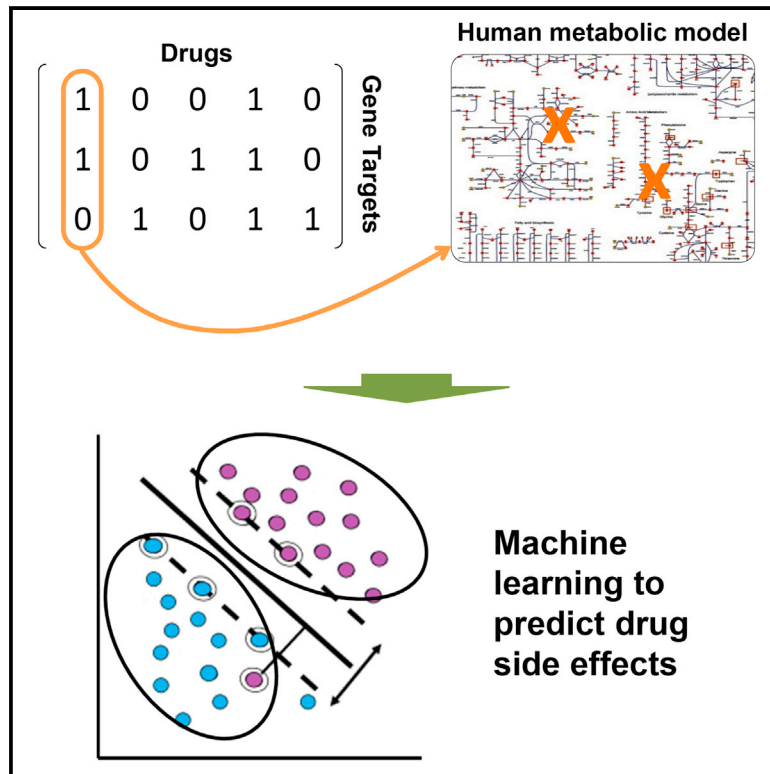


Metabolic Network Prediction of Drug Side Effects

Graphical Abstract



Authors

Itay Shaked, Matthew A. Oberhardt, Nir Atias, Roded Sharan, Eytan Ruppín

Correspondence

eyruppin@gmail.com

In Brief

A machine-learning approach is developed to predict drug side effects based on network-level signatures of metabolically acting drugs.

Highlights

- A predictor of drug side effects based on network-level metabolic signatures
- Machine learning with features derived from genome-scale metabolic modeling
- Predictors with potentially translatable significance for >70 drug side effects
- Identification of potential biomarkers and key pathways



Metabolic Network Prediction of Drug Side Effects

Itay Shaked,^{1,5} Matthew A. Oberhardt,^{1,2,3,4,5} Nir Atias,¹ Roded Sharan,¹ and Eytan Ruppin^{1,3,4,*}

¹Blavatnik School of Computer Sciences, Tel Aviv University, Tel Aviv 69978, Israel

²Department of Molecular Microbiology and Biotechnology, Faculty of Life Sciences, Tel Aviv University, Tel Aviv 69978, Israel

³Sackler School of Medicine, Tel Aviv University, Tel Aviv 69978, Israel

⁴Department of Computer Science and Center for Bioinformatics and Computational Biology, University of Maryland, College Park, MD 20742, USA

⁵Co-first author

*Correspondence: eyruppin@gmail.com

<http://dx.doi.org/10.1016/j.cels.2016.03.001>

SUMMARY

Drug side effects levy a massive cost on society through drug failures, morbidity, and mortality cases every year, and their early detection is critically important. Here, we describe the array of model-based phenotype predictors (AMPP), an approach that leverages medical informatics resources and a human genome-scale metabolic model (GSMM) to predict drug side effects. AMPP is substantially predictive (AUC > 0.7) for >70 drug side effects, including very serious ones such as interstitial nephritis and extrapyramidal disorders. We evaluate AMPP's predictive signal through cross-validation, comparison across multiple versions of a side effects database, and co-occurrence analysis of drug side effect associations in scientific abstracts (hypergeometric p value = $2.2e-40$). AMPP outperforms a previous biochemical structure-based method in predicting metabolically based side effects (aggregate AUC = 0.65 versus 0.59). Importantly, AMPP enables the identification of key metabolic reactions and biomarkers that are predictive of specific side effects. Taken together, this work lays a foundation for future detection of metabolically grounded side effects during early stages of drug development.

INTRODUCTION

Drug side effects are major causes of death (Lazarou et al., 1998), morbidity (Pirmohamed et al., 2004), and late-stage failures in drug development (Kola and Landis, 2004; Hay et al., 2014). Reliable approaches for early prediction of drug side effects are clearly needed, and they would benefit from systems analysis platforms that are broadly predictive across human cell types and tissues. An attractive system for such broad-scale analyses is cellular metabolism, which is a critical actor in many human diseases and phenotypes (Bordbar and Palsson, 2012). Metabolism is the only genome-wide network to be reliably converted into predictive models, termed genome-scale metabolic models (GSMMs) (Oberhardt et al., 2009; Thiele et al., 2013). The release of GSMMs for humans (Thiele et al., 2013) has

enabled large-scale metabolic analysis of many diseases (Lewis et al., 2010; Zelezniak et al., 2010; Shlomi et al., 2009; Duarte et al., 2007) and tissue-specific behaviors, such as the functions of liver metabolism (Shlomi et al., 2008; Jerby et al., 2010), obesity (Mardinoglu et al., 2013), and cancer (Folger et al., 2011; reviewed in Mardinoglu and Nielsen, 2012 and Bordbar and Palsson, 2012). Moreover, GSMMs provide a promising approach for predicting gene-to-phenotype linkages, including drug side effect associations.

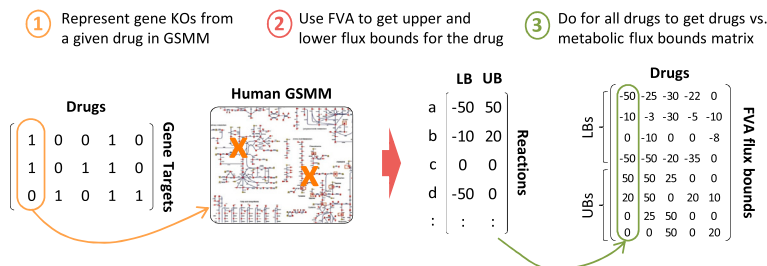
One notable recent study used GSMMs, coupled with gene expression data from human cell lines treated with a variety of drugs, to identify the molecular underpinnings of the drugs' side effects (Zielinski et al., 2015). While this is an important and timely goal, the authors did not focus on applying such knowledge in a machine-learning fashion to predict which side effects a new drug might cause, a distinct and complementary goal. Several other computational methods have emerged recently to address this need (see, e.g., Cheng et al., 2013). These methods have typically leveraged large databases of known drug side effects along with additional chemical and/or biological information and machine learning (Yamanishi et al., 2012), and some have gone further and predicted side effect-associated chemical motifs (Duran-Frigola and Aloy, 2013; Juan-Blanco et al., 2015; Pauwels et al., 2011). None, though, to our knowledge, have directly leveraged GSMMs alongside machine learning for the purpose of drug side effect prediction. Here we used machine learning to integrate bottom-up approaches (i.e., the manually curated human GSMM) with the huge reservoir of top-down data available for human disease, drugs, and their associated phenotypes into ensembles of metabolically associated side effect predictors.

RESULTS

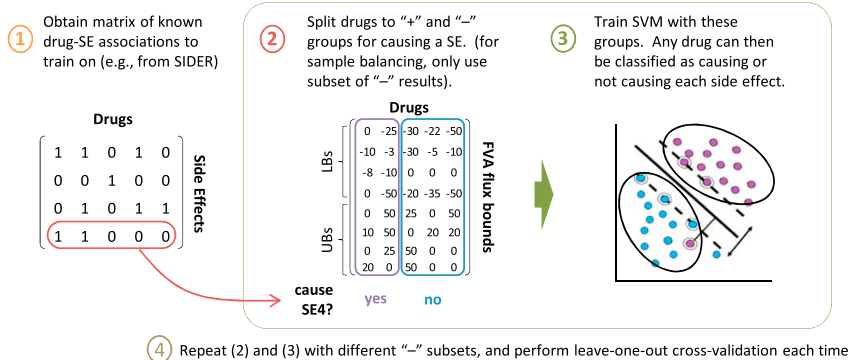
Model-Based Phenotype Predictors

We present model-based phenotype predictors (MPPs) to predict whether a side effect is caused by a given metabolically acting drug (Figure 1). MPPs are support vector machines (SVMs), which classify drugs as either causing or not causing a side effect by examining their effect on genome-wide metabolic fluxes. Each MPP predicts a single side effect, so an array of MPPs (AMPP) is constructed to predict all of the potential side effects of a given drug. An MPP for a given side effect is trained on a sample-balanced set of drugs known to and not to cause the side effect (Experimental Procedures). The features

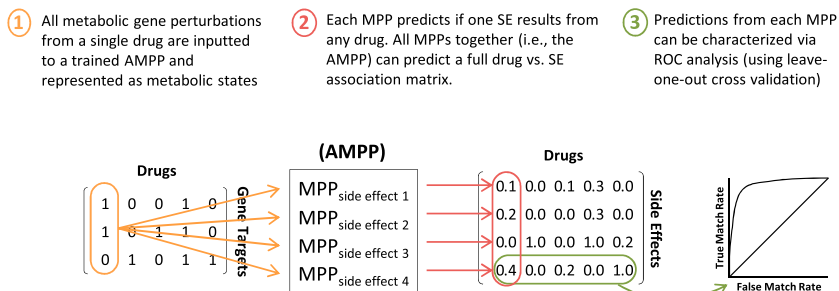
A Converting drugs to metabolic states within each MPP



B Training an MPP for a given Side Effect



C AMPP Workflow



fed into the SVM are reaction bounds for a given drug as predicted using flux variability analysis (FVA) (Mahadevan and Schilling, 2003) in a human GSMM after inactivation of the drug’s targets (Duarte et al., 2007). Because inactivation is more readily simulated in GSMMs than over-activation, only drugs that inactivate their targets are considered (this includes the large majority of all drugs). Thus, side effects for a metabolically acting drug with known targets can be predicted by first simulating the intracellular metabolic state caused by the drug (via FVA) and then feeding this state into an AMPP that predicts which side effects should occur, using a different trained SVM for each side effect.

Although we focus on drug side effects, the AMPP method can link any metabolically inactivating genetic perturbations with associated phenotypes. We test this in context of clinical symptoms and signs of inborn error of metabolism (IEM) diseases (Supplemental Experimental Procedures, Note 1; Data S1).

Figure 1. Workflow for Building an AMPP and Using It to Predict Drug Side Effects

(A) Within each MPP, inactivated drug target genes are represented as knockouts in a human GSMM. A matrix is formed by representing each drug as a column of upper and lower flux bounds as calculated via FVA (these are the features for a SVM). (B) Each MPP (one per side effect) is trained using a known drug side effect matrix. Drugs are split into those known to cause the side effect and those that don’t, and then an SVM is trained for the side effect multiple times, with sample-balanced subsets of drugs used each time. Each trained MPP is then tested using a leave-one-out cross-validation procedure. Reported AUCs are usually averages across different sample-balanced training sets. (C) The resulting AMPPs are used to predict the side effects of a given drug. First, known targets of the drug are fed into the AMPP. Each MPP within the AMPP then predicts whether inactivating the targets will cause a specific side effect. Prediction accuracy is quantified using a standard AUC/receiver operating characteristic (ROC) analysis.

Prediction of Drug Side Effects by AMPP

To build an AMPP for predicting drug side effects, we first extracted data from the side effect resource (SIDER, Kuhn et al., 2010) and DrugBank (Knox et al., 2011) and built a training set of drugs with known side effects and targets (Supplemental Experimental Procedures). We screened for the drugs that exclusively inactivate metabolic targets (since predictions are made using a metabolic model), yielding 89 metabolic-targeting drugs with known associations to 286 side effects (after filtering out side effects associated with less than ten drugs). Training an AMPP with this data yielded 217 significantly predictive MPPs (after multiple hypothesis correction), 70 of which (~25%) had areas

under the curve (AUCs) higher than 0.7 and, thus, might be translationally relevant (Figure 2A; Data S2, Average AUCs for Side Effects). The top five predictors have AUCs in the range of 0.88–0.99 and operating points at 80%–100% discovery rate with 10%–20% false positive rate.

Our strict filtering for drugs that exclusively inactivate metabolic targets is desirable, because the SVM features are based on a GSMM analysis (FVA) that best captures metabolic enzyme inactivation. However, this filtering has the disadvantage that it significantly reduces the number of drugs and side effects that can be predicted. We therefore explored how the “metabolicness” of a drug’s targets influences our predictions in two ways.

First, we rebuilt the AMPP after relaxing the requirement that all of the drug’s targets must be metabolic (Supplemental Experimental Procedures). This yielded 426 side effects associated with 190 drugs, doubling the number of metabolic pathways hit as primary targets (from 27 to 53 of 98 total in the human




Figure 2. AMPP Predictions

(A) AUC score distributions for the 37 side effects with mean AUCs greater than 0.75, as evaluated over 100 leave-one-out runs with a different (sample-balanced) set of negative training samples each time, are shown (central marks are medians, boxes are quartiles).

(B) Average AUCs correlate with the metabolism (see text) of the side effects, in that more metabolically grounded side effects are predicted better by our AMPP. Box edges are 25th and 75th percentiles and whiskers span data points not considered outliers.

(C) The distribution of AUCs across all side effect predictors is highest for AMPP-CCA, middle for AMPP, and lowest for the H-CCA method. Red dotted lines indicate mean AUC values.

(legend continued on next page)

GSMM), but keeping the distribution of pathways highly consistent (Figure S1). Of the 426 side effects, 278 (65%) were significantly predictive and 78 (18%) had AUCs higher than 0.7. We then compared the AMPP-predicted drug side effect association strengths for drug side effect pairs that have been newly added to SIDER (since the version of the database on which we trained) versus those that were newly removed. We found, reassuringly, that AMPP scored the newly added pairs significantly higher than the removed ones ($p = 2.5e-3$, rank sum; Figure 2B).

Second, to better quantify how noise from non-metabolic effects might reduce our accuracy, we assigned each side effect in the 286 highly filtered AMPP set a metabolicness score, denoting the percentage of genes that are metabolic among all genes targeted by all drugs known to cause that side effect. Reassuringly, we found that the metabolicness of a side effect correlates with the AUC of our predictor for that side effect ($\rho = 0.28$, $p = 1.2e-6$ in Pearson correlation across all 286 side effects; $\rho = 0.57$, $p = 2.1e-3$ when considering only side effects with AUCs in the top 10%; Figure 2C).

These results emphasize that the AMPP method is most predictive for metabolically targeting drugs and highly metabolically associated side effects, but that accurate predictions can still be made for many side effects even when drugs targeting some non-metabolic genes are considered.

To better understand the value of metabolic network analysis in predicting drug side effects, we compared the AMPP to a highly performing previous method (H-CCA) for predicting drug side effects, which compares biochemical structures of drugs via the machine-learning method of canonical correlation analysis (CCA) to make new drug side effect predictions (Supplemental Experimental Procedures; Atias and Sharan, 2011). We ran H-CCA on the metabolic side effects and drugs we had previously analyzed (excluding nine of the 89 drugs due to inaccessibility of their chemical structures, Supplemental Experimental Procedures), and we found that the AMPP produced a considerably higher distribution of AUCs than H-CCA across the same side effects (Figure 2D; AUCs calculated the same way for each method).

To combine the strengths of both methods, we next developed the AMPP-CCA algorithm, which combines metabolic states determined by FVA (instead of biochemical structures) with the CCA algorithm. This differs from SVM in that it integrates information from all known drug side effect interactions in predicting any one drug side effect association (Supplemental Experimental Procedures). AMPP-CCA outperformed both AMPP and H-CCA in predicting metabolic drug side effects (Figure 2D), and thus it should be considered for the task of predicting drug side effect interactions. However, in contrast to the AMPP method, AMPP-CCA cannot be drawn upon for feature selection and prediction of biomarkers (see next section), and

it requires a full drug side effect matrix to predict any one side effect.

Predictive Features and Metabolic Biomarkers for Drug Side Effects

We used a random-forest feature selection method to identify the most predictive reactions for specific side effects (Breiman, 2001; Supplemental Experimental Procedures). Notably, a handful of reactions were often sufficient to predict side effects with accuracy comparable to that of the full MPP (Figure 2E; Data S2, Side Effect Key Reactions). Purine metabolism, hyaluronan metabolism, and salvage pathways were the most enriched pathways for top predictive reactions (Data S2, Side Effect Key Pathways). Reassuringly, many of the highest confidence feature/side effect pairs were supported by previous literature (eight of the top 14 feature/side effect pairs; Data S2, Side Effect Top Reaction Set).

We next performed a second feature selection process focusing only on cellular exchange reactions, as these may point to metabolites that could be detectable in bodily fluids and thus serve as biomarkers (Experimental Procedures). This yielded 218 significantly predictive classifiers (of 286 side effects total), which we then narrowed to 162 by examining only side effects that can be predicted using between two and eight top biomarkers while attaining an AUC > 0.6 (Data S2, Side Effect Biomarkers). We found that our predicted biomarker/side effect pairs co-occur in PubMed abstracts (Shlomi et al., 2009) significantly more commonly than random pairs (Figure 2F; hypergeometric p value = $2.2e-40$). Finally, we incrementally built a predicted biomarker panel, which we increased one biomarker at a time starting with the most broadly predictive exchange metabolite, while each time assessing how many side effects the panel could predict above a significance threshold. Encouragingly, only five to ten biomarkers were sufficient for achieving an AUC > 0.6 for >100 side effects (Figures 2G and 2H). The AMPP is thus a promising method for determining small diagnostic panels of measurable biomarkers that can predict a wide array of side effects.

DISCUSSION

In summary, we used machine learning to predict phenotype-genotype relationships based on steady-state metabolic fluxes in a genome-scale model of a generic human cell. While this representation is highly simplified compared to the many temporal, spatial, and biological scales through which genotype is converted to phenotype, our approach successfully predicts gene perturbation-to-macro-phenotype associations on a subset of the side effects studied. This is even despite the fact that the approach only accounts for pharmacodynamics (i.e., drug target

(D) AMPP-based drug side effect association scores are significantly higher for associations added or removed from SIDER since the version on which we trained (to obtain sufficient sample size, drugs that do not solely target metabolism were included, along with their associated side effects).

(E) The percentage of biomarker/side effect associations co-occurring in PubMed records is higher for predicted associations than for random ones.

(F) AUCs are shown for aggregated ROC for the classification of the side effects with significant predictors, using varying numbers of features.

(G) Histogram of the number of biomarkers predictive for varying numbers of drug side effects. The biomarkers that predict the most side effects are listed in the inset box.

(H) The numbers of side effects predicted above various cutoffs, using differently sized biomarker panels, are shown.

SE, side effect.

effects) and not pharmacokinetics (i.e., effects related to clearing the drugs), as the linkages of side effects to pharmacokinetic mechanisms are generally not known. Organ specificity suggests a key area for future improvement: by using organ-specific constraint-based models (Shlomi et al., 2008; Jerby et al., 2010; Yizhak et al., 2014), we might obtain more accurate and relevant side effect predictions. Emphasis on the liver, for example, could be a key facet in future refinements of the method, since most drugs are processed in the liver and thus many side effects may be strongly associated with liver metabolism.

The AMPPs we developed can be used to predict side effects for other drugs that target metabolism, including those not yet included in our dataset. In this way, AUCs we list based on these leave-one-out validations provide important benchmarks for how well we can predict the phenotypes/side effects caused by any new drug.

SUPPLEMENTAL INFORMATION

Supplemental Information includes Supplemental Experimental Procedures, three figures, and two data files and can be found with this article online at <http://dx.doi.org/10.1016/j.cels.2016.03.001>.

AUTHOR CONTRIBUTIONS

I.S. and M.A.O. performed the analyses and wrote the paper. N.A. and R.S. did the CCA analyses. E.R. designed experiments and edited the manuscript.

ACKNOWLEDGMENTS

We thank Keren Yizhak and Allon Wagner for helpful comments. E.R. gratefully acknowledges support from the EU FP7 INFECT project, the I-CORE Program of the Planning and Budgeting Committee, the Israel Science Foundation (grant No 41/11), and CONTIBUGS EU bacterial research grants.

Received: August 3, 2015

Revised: December 11, 2015

Accepted: March 1, 2016

Published March 23, 2016

REFERENCES

- Atias, N., and Sharan, R. (2011). An algorithmic framework for predicting side effects of drugs. *J. Comput. Biol.* *18*, 207–218.
- Bordbar, A., and Palsson, B.O. (2012). Using the reconstructed genome-scale human metabolic network to study physiology and pathology. *J. Intern. Med.* *271*, 131–141.
- Breiman, L. (2001). Random forests. *Mach. Learn.* *45*, 5–32.
- Cheng, F., Li, W., Wang, X., Zhou, Y., Wu, Z., Shen, J., and Tang, Y. (2013). Adverse drug events: database construction and in silico prediction. *J. Chem. Inf. Model.* *53*, 744–752.
- Duarte, N.C., Becker, S.A., Jamshidi, N., Thiele, I., Mo, M.L., Vo, T.D., Srivas, R., and Palsson, B.O. (2007). Global reconstruction of the human metabolic network based on genomic and bibliomic data. *Proc. Natl. Acad. Sci. USA* *104*, 1777–1782.
- Duran-Frigola, M., and Aloy, P. (2013). Analysis of chemical and biological features yields mechanistic insights into drug side effects. *Chem. Biol.* *20*, 594–603.
- Folger, O., Jerby, L., Frezza, C., Gottlieb, E., Ruppin, E., and Shlomi, T. (2011). Predicting selective drug targets in cancer through metabolic networks. *Mol. Syst. Biol.* *7*, 501.
- Hay, M., Thomas, D.W., Craighead, J.L., Economides, C., and Rosenthal, J. (2014). Clinical development success rates for investigational drugs. *Nat. Biotechnol.* *32*, 40–51.
- Jerby, L., Shlomi, T., and Ruppin, E. (2010). Computational reconstruction of tissue-specific metabolic models: application to human liver metabolism. *Mol. Syst. Biol.* *6*, 401.
- Juan-Blanco, T., Duran-Frigola, M., and Aloy, P. (2015). IntSide: a web server for the chemical and biological examination of drug side effects. *Bioinformatics* *31*, 612–613.
- Knox, C., Law, V., Jewison, T., Liu, P., Ly, S., Frolkis, A., Pon, A., Banco, K., Mak, C., Neveu, V., et al. (2011). DrugBank 3.0: a comprehensive resource for 'omics' research on drugs. *Nucleic Acids Res.* *39*, D1035–D1041.
- Kola, I., and Landis, J. (2004). Can the pharmaceutical industry reduce attrition rates? *Nat. Rev. Drug Discov.* *3*, 711–715.
- Kuhn, M., Campillos, M., Letunic, I., Jensen, L.J., and Bork, P. (2010). A side effect resource to capture phenotypic effects of drugs. *Mol. Syst. Biol.* *6*, 343.
- Lazarou, J., Pomeranz, B.H., and Corey, P.N. (1998). Incidence of adverse drug reactions in hospitalized patients: a meta-analysis of prospective studies. *JAMA* *279*, 1200–1205.
- Lewis, N.E., Schramm, G., Bordbar, A., Schellenberger, J., Andersen, M.P., Cheng, J.K., Patel, N., Yee, A., Lewis, R.A., Eils, R., et al. (2010). Large-scale in silico modeling of metabolic interactions between cell types in the human brain. *Nat. Biotechnol.* *28*, 1279–1285.
- Mahadevan, R., and Schilling, C.H. (2003). The effects of alternate optimal solutions in constraint-based genome-scale metabolic models. *Metab. Eng.* *5*, 264–276.
- Mardinoglu, A., and Nielsen, J. (2012). Systems medicine and metabolic modelling. *J. Intern. Med.* *271*, 142–154.
- Mardinoglu, A., Agren, R., Kampf, C., Asplund, A., Nookaew, I., Jacobson, P., Walley, A.J., Froguel, P., Carlsson, L.M., Uhlen, M., and Nielsen, J. (2013). Integration of clinical data with a genome-scale metabolic model of the human adipocyte. *Mol. Syst. Biol.* *9*, 649.
- Oberhardt, M.A., Palsson, B.O., and Papin, J.A. (2009). Applications of genome-scale metabolic reconstructions. *Mol. Syst. Biol.* *5*, 320.
- Pauwels, E., Stoven, V., and Yamanishi, Y. (2011). Predicting drug side-effect profiles: a chemical fragment-based approach. *BMC Bioinformatics* *12*, 169.
- Pirmohamed, M., James, S., Meakin, S., Green, C., Scott, A.K., Walley, T.J., Farrar, K., Park, B.K., and Breckenridge, A.M. (2004). Adverse drug reactions as cause of admission to hospital: prospective analysis of 18 820 patients. *BMJ* *329*, 15–19.
- Shlomi, T., Cabili, M.N., Herrgård, M.J., Palsson, B.O., and Ruppin, E. (2008). Network-based prediction of human tissue-specific metabolism. *Nat. Biotechnol.* *26*, 1003–1010.
- Shlomi, T., Cabili, M.N., and Ruppin, E. (2009). Predicting metabolic biomarkers of human inborn errors of metabolism. *Mol. Syst. Biol.* *5*, 263.
- Thiele, I., Swainston, N., Fleming, R.M., Hoppe, A., Sahoo, S., Aurich, M.K., Haraldsdottir, H., Mo, M.L., Rolfsson, O., Stobbe, M.D., et al. (2013). A community-driven global reconstruction of human metabolism. *Nat. Biotechnol.* *31*, 419–425.
- Yamanishi, Y., Pauwels, E., and Kotera, M. (2012). Drug side-effect prediction based on the integration of chemical and biological spaces. *J. Chem. Inf. Model.* *52*, 3284–3292.
- Yizhak, K., Le Dévédéc, S.E., Rogkoti, V.M., Baenke, F., de Boer, V.C., Frezza, C., Schulze, A., van de Water, B., and Ruppin, E. (2014). A computational study of the Warburg effect identifies metabolic targets inhibiting cancer migration. *Mol. Syst. Biol.* *10*, 744.
- Zelezniak, A., Pers, T.H., Soares, S., Patti, M.E., and Patil, K.R. (2010). Metabolic network topology reveals transcriptional regulatory signatures of type 2 diabetes. *PLoS Comput. Biol.* *6*, e1000729.
- Zielinski, D.C., Filipp, F.V., Bordbar, A., Jensen, K., Smith, J.W., Herrgard, M.J., Mo, M.L., and Palsson, B.O. (2015). Pharmacogenomic and clinical data link non-pharmacokinetic metabolic dysregulation to drug side effect pathogenesis. *Nat. Commun.* *6*, 7101.