



Universiteit
Leiden
The Netherlands

XomicsToModel: omics data integration and generation of thermodynamically consistent metabolic models

Preciat, G.A.; Wegrzyn, A.B.; Luo, X.; Thiele, I.; Hankemeier, T.; Fleming, R.M.

Citation

Preciat, G. A., Wegrzyn, A. B., Luo, X., Thiele, I., Hankemeier, T., & Fleming, R. M. (2025). XomicsToModel: omics data integration and generation of thermodynamically consistent metabolic models. *Nature Protocols*, 21(6), 2665-2706. doi:10.1038/s41596-025-01288-9

Version: Publisher's Version

License: [Licensed under Article 25fa Copyright Act/Law \(Amendment Taverne\)](#)

Downloaded from: <https://hdl.handle.net/1887/4308606>

Note: To cite this publication please use the final published version (if applicable).

Protocol



XomicsToModel: omics data integration and generation of thermodynamically consistent metabolic models

German Preciat^{1,2}, Agnieszka B. Wegrzyn¹, Xi Luo³, Ines Thiele^{3,4,5,6}, Thomas Hankemeier¹✉ & Ronan M. T. Fleming^{1,3}✉

Abstract

Constraint-based modeling can mechanistically simulate the behavior of a biochemical system, permitting hypothesis generation, experimental design and interpretation of experimental data, with numerous applications, especially the modeling of metabolism. Given a generic model, several methods have been developed to extract a context-specific, genome-scale metabolic model by incorporating information used to identify metabolic processes and gene activities in each context. However, the existing model extraction algorithms are unable to ensure that a context-specific model is thermodynamically flux consistent. Here we introduce XomicsToModel, a semiautomated pipeline that integrates bibliomic, transcriptomic, proteomic and metabolomic data with a generic genome-scale metabolic reconstruction, or model, to extract a context-specific, genome-scale metabolic model that is stoichiometrically, thermodynamically and flux consistent. One of the key advantages of the XomicsToModel pipeline is its ability to seamlessly incorporate omics data into metabolic reconstructions, ensuring not only mechanistic accuracy but also physicochemical consistency. This functionality enables more accurate metabolic simulations and predictions across different biological contexts, enhancing its utility in diverse research fields, including systems biology, drug development and personalized medicine. The XomicsToModel pipeline is exemplified for extraction of a specific metabolic model from a generic metabolic model; it enables omics data integration and extraction of physicochemically consistent mechanistic models from any generic biochemical network. It can be implemented by anyone who has basic MATLAB programming skills and the fundamentals of constraint-based modeling.

Key points

- XomicsToModel is a semi-automated pipeline that integrates bibliomic, transcriptomic, proteomic and metabolomic data with a generic genome-scale metabolic reconstruction or model.
- It enables the seamless incorporation of multiomics datasets into metabolic reconstructions, ensuring mechanistic accuracy and physicochemical consistency.

Key references

- Luo, X., El Assal, D. C., Liu, Y., Ranjbar, S. & Fleming, R. M. T. Constraint-based modeling of bioenergetic differences between synaptic and non-synaptic components of human dopaminergic neurons. *Front. Comput. Neurosci.* **19**, 594330 (2025): <https://doi.org/10.3389/fncom.2025.1594330>
- Luo, X., Liu, Y., Balck, A., Klein, C. & Fleming, R. M. T. Identification of metabolites reproducibly associated with Parkinson's disease via meta-analysis and computational modelling. *NPJ Parkinsons Dis.* **10**, 126 (2024): <https://doi.org/10.1038/s41531-024-00732-z>
- Zagare, A. et al. Omics data integration suggests a potential idiopathic Parkinson's disease signature. *Commun. Biol.* **6**, 1179 (2023): <https://doi.org/10.1038/s42003-023-05548-w>

¹Metabolomics and Analytics Center, Leiden Academic Center for Drug Research, Leiden University, Einsteinweg, Leiden, The Netherlands. ²Department of Translational Bioengineering, University of Guadalajara, Guadalajara, Mexico. ³School of Medicine, University of Galway, Galway, Ireland. ⁴Ryan Institute, University of Galway, Galway, Ireland. ⁵Division of Microbiology, University of Galway, Galway, Ireland. ⁶APC Microbiome Ireland, Galway, Ireland. ✉e-mail: hankemeier@lacdr.leidenuniv.nl; ronan.mt.fleming@gmail.com

Introduction

Genome-scale metabolic models

Genome-scale metabolic models are computational frameworks that simulate the metabolic processes within an organism. They are essential tools for understanding complex biochemical networks and predicting physiochemical and biochemically realistic metabolic fluxes. The main goal of a genome-scale metabolic model is to represent known metabolic functions in living systems to facilitate hypotheses generation, experimental design and the interpretation of experimental data¹.

Genome-scale metabolic models can be generic, covering multiple tissues, organisms or organelles, or they can be specific, tailored to specific tissues, cell types, or environmental conditions. There are four complementary approaches for the development of genome-scale metabolic models:

1. From scratch: if no network reconstruction is available, a high-quality genome-scale metabolic reconstruction of a single organism can be generated from scratch by following established protocols². Once this reconstruction is in place, it can be used to derive various metabolic models by the computational application of different combinations of mathematical modeling assumptions.
2. Generic reaction databases: given a generic metabolic reaction database, containing metabolic reactions from multiple organisms, an organism-specific model can be derived directly from these data. There are several existing methods for doing this, for example, CarveMe³.
3. Orthologous genes: if a network reconstruction, or a model, is available for an organism with orthologous genes, it is possible to generate metabolic models for related species. Methods for microbial⁴ and mammalian⁵ species have been proposed.
4. Tailoring models: given a generic metabolic reconstruction or model for a multicellular organism (containing metabolic reactions from any anatomical region), there are several methods that can be used to derive a model that is specific for a particular anatomical region⁶.

In all except the first instance, the experiment starts with an existing generic reconstruction or model, and the method extracts a specific model, which is a subset of the generic reconstruction.

A specific model may be specific to a particular organism or a particular anatomical region. For example, given a generic human metabolic model, one could extract a hepatocyte model⁷, with the capability to model different contexts, depending on the constraints subsequently applied to it. A specific model may also be context-specific, which is a particular organism or anatomical region in a particular environmental and internal context. From the same starting point, with different input data, one could extract a context-specific model, for example, a hepatocyte model that is specific to a fasting state⁸. Context-specific data can include both experimental results and bibliomic data. When discrepancies arise, such as genes, metabolites, or reactions present in one data source but absent in another, preference must be given to either experimental or bibliomic data, resulting in two types of model: cell-type-specific and condition-specific models. Cell-type-specific models prioritize experimental data, whereas condition-specific models prioritize bibliomic data^{9,10}.

To ensure the accuracy and biological relevance of a context-specific model, it is crucial to establish stoichiometric, flux and thermodynamic consistency. Stoichiometric consistency guarantees that all reactions within the model are mass-balanced, ensuring that the inputs and outputs of each reaction align with the conservation of mass¹¹. Flux consistency ensures that every reaction in the model can carry a nonzero steady-state flux, facilitating meaningful metabolic simulations¹². Lastly, thermodynamic flux consistency ensures that the model adheres to the laws of thermodynamics, meaning that each reaction can carry a nonzero net flux in a thermodynamically feasible direction, respecting energy conservation and equilibrium constraints¹³. These consistency checks are essential to generate a physiologically plausible and predictive metabolic model.

Established model extraction methods ensure that a context-specific model is flux consistent. This concept is explained mathematically after a brief definition of the mathematical notation used throughout.

- $\mathbb{R}$ denotes the field of real numbers
- $\mathbb{R}_n$ denotes the vector space of n -tuples of real numbers
- $\mathbb{R}^{m \times n}$ denotes the space of $m \times n$ matrices with entries in $\mathbb{R}$
- Greek lowercase and uppercase letters denote scalars and sets, respectively
- Roman lowercase and uppercase letters denote vectors and matrices, respectively
- N^T denotes the transpose of a matrix N
- $\|r\|_0$ denotes the zero norm of the vector r

In any optimization problem, the variables are indicated below the first symbol in the optimization problem definition, for example, $\max_{v \in \mathbb{R}^r}$ denotes the maximization of an objective in an optimization problem where the only variable is v and all other symbols are input data to the optimization problem.

To predict the steady-state behavior of metabolic networks, flux balance analysis is commonly employed as it provides an efficient and widely recognized method for calculating metabolic fluxes under given constraints. Flux balance analysis ensures that the metabolic model adheres to mass balance and other physiological constraints, making it a foundational approach for exploring metabolic capabilities. This method serves as a stepping stone for more advanced flux analysis approaches, which may offer additional insights or higher complexity, but flux balance analysis remains essential for establishing baseline flux distributions.

Flux balance analysis

Flux balance analysis¹⁴ requires the solution to the following optimization problem:

$$\begin{aligned} \max_{v \in \mathbb{R}^r} \rho(v) &:= c^T v \\ \text{s.t. } Sv &= 0 \\ l &\leq v \leq u, \end{aligned} \quad (1)$$

where $v \in \mathbb{R}^r$ is the net rate of each reaction, and $\rho(v) := c^T v$ is a biologically motivated linear objective, specified by the coefficient vector $c \in \mathbb{R}^r$. The matrix $S \in \mathbb{R}^{m \times r}$ is a stoichiometric matrix, where m is the number of metabolites, and r is the number of internal plus exchange reactions. The constraint $Sv = 0$ represents the assumption of steady state, that is:

- Production = consumption for metabolites not exchanged across the boundary of the system
- Production + uptake = consumption + secretion for metabolites exchanged across the boundary of the system

The inequalities $l \leq v \leq u$ denote box constraints from the lower ($l \in \mathbb{R}^r$) and upper bounds ($u \in \mathbb{R}^r$) on reaction rates.

A reaction is considered flux consistent if it can carry a nonzero steady-state flux, meaning there is at least one steady-state flux vector v in the feasible set $\Omega := \{Sv = 0, l \leq v \leq u\}$ where $v_j \neq 0$. This implies the reaction can function under the given conditions and constraints. Conversely, a reaction is considered flux inconsistent if its rate $v_j = 0$ for all steady state flux vectors in the feasible set Ω . As flux balance analysis only predicts steady-state flux, including flux-inconsistent reactions in a model can be misleading. Therefore, it is important to distinguish between:

- Reactions that cannot carry steady-state flux (independent of the objective chosen) and
- Reactions that do not carry steady-state flux in a particular flux balance analysis solution (dependent of the objective chosen)

Therefore, flux-inconsistent reactions should be identified in a model before interpreting the results of methods such as flux balance analysis. In a metabolic reconstruction, a flux-inconsistent reaction (supported by experimental evidence) may serve as a starting point for refinement (for example, by gap filling¹⁵) that adds one or more additional reactions to a network to enable flux through a formerly flux-inconsistent reaction.

Thermodynamic feasibility and thermodynamic flux consistency

It is well established that flux balance analysis requires additional constraints to ensure that a steady-state flux vector is also thermodynamically feasible^{16,17}, that is, consistent with energy conservation and the second law of thermodynamics¹⁸. The critical necessity at genome-scale is to incorporate any additional constraints (or terms in the objective) in a mathematical form that, when computationally implemented, retains the computational efficiency and certificates available with linear optimization. As exemplified in Fig. 1, a stoichiometric matrix may be split into two sets of columns $S := [N, B]$, where the matrix $N \in \mathbb{R}^{m \times n}$ represents internal reactions, which do conserve mass, and the matrix $B \in \mathbb{R}^{m \times k}$ represents external reactions, which do not conserve mass, for example, the reaction with equation $A \rightarrow \emptyset$. Within this framework, chemical reactions that do not conserve mass serve as modeling constructs for simulating the exchange, demand or sink of metabolites with the environment. Exchange reactions model the flux of metabolites into or out of the system, simulating substance exchange with the external environment. Demand reactions represent the cellular need for a specific metabolite, mimicking the consumption or requirement within the biological system. Lastly, sink reactions act as outlets for metabolites, replicating their removal from the system, whether through excretion or degradation or being otherwise removed from the cellular environment.

By splitting internal net fluxes into unidirectional fluxes, and maximizing the entropy of unidirectional fluxes, one can compute a thermodynamically feasible steady-state net flux vector in a genome-scale biochemical model¹³. Maximizing the entropy of unidirectional fluxes is a convex optimization problem that retains the efficiency as well as feasibility certificates available with linear optimization¹⁹. This approach of maximizing the entropy of unidirectional fluxes assumes that the model constraints admit a solution that corresponds to a thermodynamically feasible net flux vector. However, we define a model to be thermodynamically flux consistent if each of its reactions admits a nonzero net flux in at least one thermodynamically feasible, steady-state, net flux vector. Chemically, a thermodynamically feasible net flux vector may be zero for a reaction if the substrates and products of a reaction are in thermodynamic equilibrium. Biochemically, a thermodynamically feasible net flux vector may be zero for a reaction (effectively zero at a biochemically relevant timescale) if the enzyme corresponding to that reaction is absent, even if the substrates and products for that reaction are not in thermodynamic equilibrium. We define a reaction to be thermodynamically flux-consistent if it admits a nonzero net flux that is also thermodynamically feasible to eliminate thermodynamically inconsistent reactions that only admit a zero net flux in a thermodynamically feasible flux vector. In summary, a reaction is flux consistent if it admits a nonzero flux that satisfies mass balance and bound constraints, whereas it is thermodynamically flux consistent if it admits a nonzero net flux that satisfies a combination of mass balance, bound constraints, energy conservation and the second law of thermodynamics.

A thermodynamically feasible flux does not exist when at least one combination of bounds on reaction rates forces net flux around a stoichiometrically balanced cycle²⁰. Consider the




Fig. 1 | Network diagram of upper glycolysis with reactions and metabolites with corresponding labels. Left: a network diagram of upper glycolysis (left) with reactions (arrows, uppercase labels) and metabolites (nodes, lowercase labels) with corresponding labels. Right: the corresponding stoichiometric matrix ($S \in \mathbb{R}^{13 \times 10}$, $N \in \mathbb{R}^{13 \times 8}$, $B \in \mathbb{R}^{13 \times 2}$), with reversible reactions shown in green, nonreversible reactions shown in blue and exchange reactions shown in red. The upper bound (ub) of internal reactions is unconstrained, whereas the lower bound (lb) is limited by reaction directionality or by the maximum uptake rate, which can be seen as a constraint on uptake of extracellular glucose (glc-D[e]) or secretion of dihydroxyacetone phosphate (dhap[e]) from the environment. All metabolite and reaction abbreviations are with respect to the namespace in <https://www.vmh.life/> (ref. 55).

following stoichiometrically balanced cycle of reactions $A \rightleftharpoons B \rightleftharpoons C \rightleftharpoons A$. If the bounds are such that $A \rightarrow B \rightarrow C \rightarrow A$ is the only feasible net flux, then the network containing this cycle cannot carry a thermodynamically feasible steady state net flux. The second law of thermodynamics requires that a chemical potential difference between substrates and products is required for net flux. Let μ_a denote the chemical potential of metabolite A . Given net flux $A \rightarrow B$, the second law of thermodynamics requires $\mu_A > \mu_B$ and similarly $\mu_B > \mu_C$ and $\mu_C > \mu_A$. However, the first pair of inequalities implies $\mu_A > \mu_C$, which is inconsistent with the last, unless one assigns more than one chemical potential to C at the same instant. However, to do so would be inconsistent with energy conservation, which requires that each row of a stoichiometric matrix corresponds with a single chemical potential, assuming each row represents a metabolite in a specific compartment¹⁸.

Stoichiometrically balanced cycles are biochemically faithful network topological features that are omnipresent in genome-scale models. It is not the presence of a stoichiometrically balanced cycle, per se, that presents a problem. Thermodynamically inconsistent specification of reaction bounds may force net flux around a stoichiometrically balanced cycle and thereby prevent the prediction of a thermodynamically feasible net flux. In addition, failure to implement thermodynamic constraints on the optimal solution of an optimization problem may lead to the prediction of a thermodynamically infeasible net flux. As recognized by Desouki et al.²⁰, provided that a model does not contain a combination of thermodynamically inconsistent bounds, the minimization of the one-norm of net flux, subject to additional constraints maintaining the direction of net flux, is guaranteed to remove that part of a steady-state flux vector that is thermodynamically infeasible. This approach is attractive, as it is based on linear optimization, but it has yet to be leveraged for the generation of thermodynamically flux-consistent models. Generation of thermodynamically flux-consistent models should, in principle, increase the biochemical fidelity of constraint-based models as it opens up the possibility of leveraging established methods to efficiently compute thermodynamically feasible steady state fluxes^{13,20} that assume an input model is thermodynamically flux consistent.

Given a generic reconstruction or model, all established model extraction algorithms, including GIMME²¹, iMAT²², MBA⁷, mCADRE²³, FastCore¹² and CarveMe³, as well as gap-filling algorithms¹⁵, extract a specific model on the basis of either a binary (present/absent) or weighted assignment of reactions desired to be present in a specific model, where each reaction is (net) flux consistent. However, these algorithms do not ensure that the resulting models are thermodynamically flux consistent. This Protocol introduces an enhanced model extraction algorithm (thermoKernel algorithm) within a model generation pipeline that generates models that are thermodynamically flux consistent.

Development of the Protocol

This Protocol presents the XomicsToModel pipeline, which enables the extraction of a context-specific, genome-scale metabolic model that is stoichiometrically, thermodynamically and flux consistent from generic reconstruction and omics data. The pipeline was initially developed and validated in the context of dopaminergic neuronal metabolism⁹, where transcriptomic, metabolomic and bibliomic data were integrated with Recon3D²⁴. Beyond this first application, XomicsToModel has already been successfully used in several peer-reviewed studies. For example, it has been applied to construct context-specific genome-scale models of human dopaminergic neurons by integrating transcriptomic and exometabolomic data to investigate neuronal metabolism and its bioenergetic vulnerabilities²⁵; to integrate multiomics datasets for the derivation of condition-specific human metabolic models, demonstrating reproducibility across biomedical contexts²⁶; and to explore disease mechanisms by combining transcriptomic, proteomic and metabolomic data²⁷. In parallel, the pipeline has also been employed to predict fluxes when combined with stable isotope labeling data^{28,29}. Together, these applications demonstrate the robustness and adaptability of the pipeline across neuroscience and broader biomedical research.

Developing a comprehensive pipeline for generating thermodynamically flux consistent models is essential for enhancing the accuracy of constraint-based metabolic models. By incorporating methods such as the thermoKernel algorithm, this pipeline addresses

the limitations of existing model extraction algorithms and ensures that thermodynamic consistency is integral to model generation. This holistic approach enables more reliable predictions of metabolic behavior and flux distributions, ultimately advancing our understanding of cellular processes.

Overview of the procedure

A metabolic reconstruction is a representation of a metabolic network whose content may not entirely satisfy modeling assumptions. By contrast, a metabolic model is a mathematically defined representation that adheres to such assumptions. *XomicsToModel* is a semi-automated computational pipeline designed to generate context-specific metabolic models from a generic reconstruction or model, combined with omics data and user-defined technical parameters.

The pipeline consists of 23 sequential steps (Fig. 2), organized into four functional stages: (1) input preparation, in which omics datasets, technical parameters and reconstruction metadata are specified; (2) data integration and consistency checks, where context-specific information is mapped onto the generic reconstruction; (3) model tailoring, in which constraints are iteratively applied and internal variables refined; and (4) model finalization, where a comprehensive record of the reconstruction process is generated, including the reactions added or removed at each step. Each step is controlled by user-defined variables and, except for the final stage, produces intermediate outputs that guide iterative refinement.

Although a draft model can be generated using default parameters, optimal results require adjusting parameters to the biological context and validating model predictions against experimental data. A practical strategy is to generate an ensemble of draft models under different parameterizations, quantify their predictive performance and select the context-specific model that best aligns with experimental evidence⁹.

For clarity, the full sequence of the 23 modular steps is provided below, ensuring reproducibility and highlighting the adaptability of the pipeline to available data.

The workflow of *XomicsToModel* proceeds in a sequence of modular steps, ensuring that the extraction can adapt to the available information. It begins with data preparation (Step 1) followed by a generic model check (Step 2) to assess consistency, after which an objective function can be set if desired (Step 3). The model can then be enriched with missing reactions




Fig. 2 | Overview of the *XomicsToModel* pipeline. It is possible to create genome-scale models iteratively and systematically using the *XomicsToModel* pipeline. The generated models can be used to design or interpret the information on the basis of bibliomic data from manual literature curation (red), as well as metabolomic (blue), proteomic (green) or transcriptomic (purple) data. Furthermore, experimental data can be used to validate or refine the context-specific model. The *XomicsToModel* pipeline (gray) extracts a new model by integrating a generic model with context-specific information based on the technical parameters such as the extraction algorithm, method for identifying active genes or maximum reaction rates. Steps 12, 13 and 22 are indented as they represent subprocesses within the pipeline.

(Step 4); core metabolites and reactions are identified (Step 5); and reaction bounds are established (Step 6). Gene activity is incorporated (Step 7), with optional closure of ionic fluxes (Step 8), exchange reactions (Step 9) and sink or demand reactions (Step 10). At this stage, metabolic constraints can be introduced (Step 11), drawing from cell culture data (Step 12) and metabolomics (Step 13), while allowing custom user-defined constraints (Step 14) or coupled reactions (Step 15). The model is refined by removing inactive reactions (Step 16) or genes (Step 17), and subsets that are flux-consistent (Step 18) or thermodynamically consistent (Step 19) are identified if required. Active reactions based on gene expression are then integrated (Step 20) before the final extraction of the context-specific model (Step 21), which can also be performed using `thermoKernel` if selected (Step 22). Lastly, the workflow is completed with final adjustments and reporting (Step 23). Importantly, the process is modular: if transcriptomic, metabolomic or bibliomic data are missing, the extraction still proceeds seamlessly with the available information.

Input data

The input data to the `XomicsToModel` pipeline are a generic biochemical network and a set of omics data. The generic biochemical network may be a reconstruction or a model without any specificity and is not required to be stoichiometrically, thermodynamically or flux consistent. The `XomicsToModel` pipeline allows a flexible and modular integration of transcriptomic, proteomic and metabolomic data, as well as bibliomic data abstracted from literature curation. In each case, the input omics data may be qualitative (present, absent and unspecified), semiquantitative, quantitative or combinations thereof. The `XomicsToModel` pipeline is complemented by functions to automatically import omics data. In the `XomicsToModel` pipeline, metabolomic data on intracellular metabolites can currently be integrated only in a qualitative or semiquantitative rather than fully quantitative manner. As `XomicsToModel` enables the generation of a stoichiometrically, thermodynamically and flux-consistent model, this model is compatible with subsequent integration of quantitative intracellular metabolomic data using a variety of established techniques³⁰.

The application of the `XomicsToModel` pipeline to extract a dopaminergic neuronal metabolic model⁹ demonstrates its flexibility with respect to the incorporation of a variety of qualitative and quantitative constraints. For example, information about the presence or absence of metabolites in the culture medium, together with quantitative metabolite exchange reaction rates, was integrated into the model. These exchange rates were used to set bounds on the corresponding reactions through quadratic optimization, with each constraint weighted by the inverse of its measurement uncertainty. In the `XomicsToModel` pipeline, it is possible to manually specify metabolic reactions to extend an input reconstruction or model (with manually specified metabolic reactions) when a generic model is missing certain key pathways relevant for a system of interest.

Constraint relaxation

The `XomicsToModel` pipeline incorporates a series of tests for flux and optionally thermodynamic consistency after each step, for example, after the removal of generic metabolites or reactions assigned not to be present in a specific model. This approach ensures the early detection of infeasible constraints resulting from inaccurate or inconsistent experimental data. The detection of inconsistency is followed by the algorithmic relaxation of constraints (for example, bounds) to render the draft model feasible. Specifically, the `XomicsToModel` pipeline automatically searches for the minimal number of constraint relaxations required to admit a flux consistent model in the case that inconsistent or incorrect omics data render a draft model infeasible. This approach can be used to feed back into data processing or input to avoid model infeasibility, which is often a challenge with omics data integration. For example, the relaxation of a steady-state constraint points to the need for additional reactions to enable either the synthesis or degradation of a metabolite where the existing total rate of synthesis (production) cannot match the total rate of degradation (consumption). In this case, additional reactions may need to be added to the generic model based on following established manual reconstruction protocols² or by recourse to algorithmic gap filling¹⁵.

The application of the `XomicsToModel` pipeline to extract a dopaminergic neuronal metabolic model⁹ presented a novel challenge for constraint-based modeling, because the extracted model was required to be representative of neurons that do not grow. Therefore, the maximization of biomass growth could not be used as an objective. Instead, a set of coupling constraints³¹ was added to represent cell maintenance requirements, for example, for the turnover of key metabolites. However, if incorrectly scaled with respect to constraints on exchange reactions, the application of coupling constraints may generate an infeasible draft model. Therefore, an algorithmic proposition of a minimal number of constraint relaxations accelerates the identification of biochemically inconsistent constraints and the refinement of input data. That is, proposed relaxations should be reviewed manually for consistency with additional biochemical knowledge then adjustments made to ensure that the input data, the input generic model or both are consistent with a feasible model.

Model extraction

The `XomicsToModel` pipeline is compatible with various model extraction algorithms (Table 1), with an established interface to `FastCore`¹². In addition, the `XomicsToModel` pipeline builds upon `thermoKernel`, a novel algorithm used to extract the dopaminergic neuronal metabolic model⁹ but is applicable for the extraction of any context-specific model that is required to be stoichiometrically, thermodynamically and flux consistent. The ability of `thermoKernel` to ensure thermodynamic flux consistency during the model extraction process opens new possibilities for data integration. Thermodynamic consistency is implemented by constraining the possible relationships between reaction flux and metabolite chemical potential. Consequently, it is possible to directly specify a qualitative (presence, absence and unspecified) or quantitative (weighted) assignment of metabolites desired to be present in a specific model.

More generally, it is possible to specify any combination of qualitative assignments and quantitative weights to metabolites and reactions to reflect what is desired to be present or absent from a specific model. These weights are determined objectively based on omics evidence to avoid any subjective errors in assignment. This enables the `thermoKernel` algorithm to extract a thermodynamically flux-consistent model that is a tradeoff between qualitative assignments and/or quantitative weights of rows and columns of a stoichiometric matrix. The application of the `XomicsToModel` pipeline to extract a dopaminergic neuronal metabolic model⁹ demonstrates this flexibility with the qualitative and weighted specification of metabolites and reactions desired to be present in the dopaminergic neuronal model. For example, the `XomicsToModel` pipeline enables the integration of quantitative weights

Table 1 | Comparison of technical features of algorithms for extraction of a context-specific model, given a generic model of a single organism

Properties\algorithms	GIMME	iMAT	MBA	tINIT	mCADRE	FastCore	thermoKernel
Citation	21	22	7	53	23	12	N/A
Year published	2008	2010	2010	2014	2012	2014	2023
Metabolite list	x	x	x	x	x	x	✓
Metabolite weights $\in \mathbb{R}$	x	x	x	x	✓	x	✓
Active reaction list	✓	✓	✓	✓	✓	✓	✓
Reaction weights $\in \mathbb{R}$	x	x	✓	✓	✓	x	✓
Coupling constraints	x	x	x	x	x	✓*	✓
Thermodynamic consistency	x	x	x	x	x	x	✓
Constraint relaxation	x	x	x	x	x	x	✓
Algorithmic approach	MILP	MILP	FVA	MILP	fastFVA	QC-LP	DCA-LP

*Denotes subsequently implemented in the COBRA Toolbox³⁴. In a context-specific metabolic network, a metabolite is considered to be present if the sum of its rate of production plus consumption is above a minimum magnitude considered nonzero (typically 1×10^{-6}); otherwise, it is considered to be absent. MILP, mixed integer linear programming; FVA, flux variability analysis, fastFVA⁵⁴, computational acceleration of flux variability analysis; QC-LP, quasiconcave sequence of linear programs; DCA-LP, difference of convex function sequence of linear programs.

on reactions from transcriptomic data at the same time as qualitative assignments from metabolomic data, where the metabolite presence (absence or unknown) can now be applied directly as context-specific data input.

Ensemble modeling

The `XomicsToModel` pipeline is flexible, and various additional functions may be incorporated into it. The supplementary `XomicsToMultipleModels` function can be used to generate an ensemble of context-specific genome-scale models by varying the specific data and technical parameters, as well as experimental data or generic models used in the model extraction process. Moreover, a supplementary `modelPredictiveCapacity` function can be added to estimate the predictive capacity of an extracted model, given independent data. The creation of multiple models with varying parameters and comparison of the predictive capacity of each model enables the identification of optimal input data and technical parameters. For example, where multiple different data types are available to generate a context-specific model (for example, transcriptomic data corresponding to genes that encode metabolic enzymes and information manually curated from literature on the presence or absence of activity on particular metabolic reactions), it is not always obvious what the relative merits of those input data are, especially as they may be conflicting. For example, gene expression might indicate that the gene for an enzyme is expressed, but manual curation indicates that the corresponding reaction is not active. In this situation, a pragmatic approach is to generate an ensemble of models, compare their predictive capacity and then select the combination of input data and technical parameters that provides the most accurate prediction. Further information on ensemble modeling may be found in published literature³².

Applications

One of the first applications of `XomicsToModel` was to construct dopaminergic neuronal models, providing insight into metabolic dysfunction in Parkinson's disease⁹. Substantia nigra dopaminergic neurons are the most susceptible to degeneration in Parkinson's disease, yet their biochemical mechanisms remain poorly understood³³. This foundational work has been extended in other biomedical contexts. For example, `XomicsToModel` was employed in studies of human neuronal metabolism²⁵, in condition-specific human metabolic reconstructions integrating diverse omics datasets²⁶ and in disease-focused analyses where transcriptomic, proteomic and metabolomic integration revealed mechanistic alterations²⁷. In addition, ongoing studies have successfully used the pipeline to integrate isotopic labeling experiments with constraint-based modeling^{28,29}.

The `XomicsToModel` pipeline offers broad applications for investigating metabolic dysfunction in various diseases. It can be applied to any situation where one has a generic (or universal) reconstruction (or model) and seeks to extract a specific subset of it, based on transcriptomic, proteomic or metabolomic data or combinations thereof. This pipeline does not impose any requirements regarding the species identity in the input reconstruction or model nor does it limit the number of species represented in either the input or output models. The most computationally intensive steps are the identification of the largest thermodynamically consistent subset of an input reconstruction or model and the extraction of a thermodynamically consistent subset of minimal size, both of which are achieved with the `thermoKernel` algorithm. As the `thermoKernel` algorithm is based on a sequence of linear optimization problems, its performance scales accordingly. The qualitative specification of metabolite (presence/absence/unspecified), reaction (activity/inactivity/unspecified), gene (activity/inactivity/unspecified) and combinations thereof are all possible. Weighted specification to bias for or against the inclusion of metabolites, reactions or both (based on input transcriptomic, proteomic or metabolomic data) can be performed alone or in combination with a qualitative specification. The `XomicsToModel` pipeline is interfaced via a function, but multiple options can be specified to enable modular use of the `XomicsToModel` pipeline to implement model extraction priorities reflective of a wide variety of scenarios. As such, the `XomicsToModel` pipeline will probably have widespread applicability.

Comparison with other methods

The `XomicsToModel` pipeline is integrated within the COBRA Toolbox³⁴. It consists of a suite of functions for context-specific input data preprocessing, identification of the subset of a reconstruction that satisfies certain modeling assumptions, integration of constraints derived from context-specific data with a generic metabolic model, recovery of model feasibility if the constraints are inconsistent, extraction of a context-specific model or an ensemble thereof. Moreover, it contains a wide variety of functions for exploring the properties of generated models, including the evaluation of their predictive fidelity^{9,35}, as detailed in the Supplementary Information. Although other established pipelines^{36,37} exist that offer similar functionality, the `XomicsToModel` pipeline provides several novel capabilities, especially its ability to implement additional constraints, integrate additional data types and relax an infeasible model to recover from the application of inconsistent constraints. These features are primarily enabled by the `thermoKernel` model extraction algorithm.

The mathematical similarities and differences between established model extraction algorithms have been extensively reviewed³⁸. Furthermore, a systematic comparative evaluation of model extraction methods⁶ has shown that using various data types, such as omics data, for model construction and validation, as well as the careful selection of gene expression cutoffs, lead to higher model accuracy when compared with experimental metabolomic data. Subsequently, guidelines for extracting biologically relevant context-specific metabolic models using gene expression data were proposed³⁹. The `XomicsToModel` pipeline is implemented in a manner consistent with those guidelines, but the `thermoKernel` model extraction algorithm also adds additional functionality. Table 1 compares `thermoKernel` with other algorithms for the extraction of a context-specific model. By default, the `XomicsToModel` pipeline uses `thermoKernel`, but it has been implemented to interface to any of the model extraction algorithms in Table 1 via the generic interface provided by `createTissueSpecificModel` in the COBRA Toolbox³⁴. Currently, the `XomicsToModel` pipeline has also been extensively tested with `FastCore`¹² but could readily be interfaced with any other model extraction algorithm.

Experimental design

Required input data

To run the `XomicsToModel` pipeline, users must first obtain and prepare a generic reconstruction or model and the context-specific data used to extract the model. Furthermore, one must either accept the default set of technical parameters or specify the parameters to suit a particular context (as detailed below).

Generic reconstruction or model: a generic genome-scale reconstruction or model represents a metabolic network constructed from the amalgamation and manual literature curation of metabolic reactions that occur in various cell types or organisms. For example, a comprehensive generic reconstruction of human metabolism, such as Recon3D²⁴, which serves as a valuable resource for metabolic modeling, can be easily accessed and downloaded from the [Virtual Metabolic Human](#) database.

Context-specific data: context-specific data represent the genotype or phenotype of a specific biological system. It can be obtained through a review of existing literature or derived experimentally from a biological system or both. The information can be entered manually or using the function `preprocessingOmicsModel` described in the Supplementary Information. The following context-specific information is currently supported by the `XomicsToModel` pipeline:

- Bibliomic data are data derived from a manual reconstruction following a review of existing literature. This includes data on the genes, reactions or metabolites known to be present or absent in the studied biological system. Furthermore, bibliomics data can define a set of coupled reactions or a set of constraints for model reactions based on phenotypic observations
- Transcriptomic data are measured gene expression levels in the studied biological system. It is used to estimate qualitatively and quantitatively the activity of reactions associated with the detected genes based on the gene–protein–reaction association (Fig. 3). Transcriptomic data can be provided in fragments per kilobase million (FPKM) or raw counts




Fig. 3 | The gene–protein–reaction association are Boolean operators that describe the interactions of genes, transcripts, proteins and reactions. For the reaction glucose-6-phosphate isomerase (PGI) to occur, the gene PGI must be activated (Entrez identifier: 2821), leading to the transcription of the PGI mRNA (NCBI reference sequence: NM_000175.5). The mRNA is then translated into the glycolytic enzyme PGI (UniProt identifier: P06744). Finally, the PGI enzyme (Enzyme Commission number: 5.3.1.9) catalyzes a reaction that interconverts glucose-6-phosphate (VMH identifier: g6p) and fructose-6-phosphate (VMH identifier: f6p).

- Proteomic data are measured protein levels in the studied biological system. Similarly to transcriptomics data, they can be used to quantitatively estimate which reactions in the metabolic model should be considered active based on the gene–protein–reaction association (Fig. 3)
- Metabolomic data are the average and standard deviation of metabolite concentrations measured in cell media, biofluids, tissues or organisms translated into flux units, which may be mmol/gDW/h, μ mol/gDW/h or nmol/gDW/h, depending on personal preference. Metabolites detected experimentally can be assigned to be present in the metabolic model. Measured uptake and secretion rates in growth media or biofluids can be used to quantitatively constrain the uptakes and secretions of the model. Furthermore, growth condition information such as growth media composition can also be provided to constrain available uptakes in the model

Although *XomicsToModel* supports the input of all of these different types of context-specific information simultaneously, you only need a subset of these different types of context-specific information to generate a specific model. For example, if there are no proteomic data, it is still possible to generate a context-specific model.

Technical parameters: with these options, technical constraints can be added to the model, as well as setting the parameters for model extraction or debugging. If they are not specified, default values are used.

- Bounds are parameters that define the minimum and maximum flux value in the model or the minimum nonzero flux value tolerance
- Exchange reactions are instructions to close or leave open demand and sink reactions and whether exchange reactions should be added on the basis of the metabolomic data
- Extraction options are technical parameters that serve as inputs for solvers and extraction algorithms, which control the process of generating a context-specific model
- Data-specific parameters are parameters that define the minimum level of transcript and/or protein to be considered as present in the model (thresholds) and whether the transcripts below the defined threshold should be removed from the generic reconstruction or model. They also determine whether literature-curated data should be prioritized over other omics data and whether weights from transcriptomic data should be used when including reactions with *thermoKernel*
- Debugging options are debugging parameters that allow the user to evaluate the output of consecutive steps within the pipeline after they have been completed

Thermodynamically consistent fluxes (Step 19)

Fluxes must not only satisfy stoichiometric and steady-state constraints but also adhere to the laws of thermodynamics. In this context, a flux is considered thermodynamically feasible if its direction is consistent with a negative Gibbs free energy change, that is, flux proceeds from higher to lower chemical potential.

A reaction is defined as thermodynamically flux-consistent if it admits at least one nonzero flux vector that satisfies both stoichiometric mass balance and thermodynamic feasibility. Conversely, reactions that can only carry zero flux under all feasible thermodynamic states are classified as thermodynamically inconsistent.

To approximate the largest thermodynamically consistent subset of reactions, `XomicsToModel` employs an iterative strategy that accumulates reactions supporting feasible flux vectors until no further consistent reactions can be identified. This pragmatic approach balances computational tractability with biochemical fidelity.

In practice, Step 19 ensures that only reactions capable of carrying nonzero, thermodynamically feasible fluxes are retained for the subsequent model extraction.

Any stoichiometric matrix S may be split into one subset of columns corresponding to internal and external reactions, $S = [N, B]$, where internal reactions are stoichiometrically consistent, that is $\exists \ell \in \mathbb{R}_{\geq 0}^m$ such that $N^T \ell = 0$, and external reactions are not stoichiometrically consistent, that is $\nexists \ell \in \mathbb{R}_{\geq 0}^m$ such that $B^T \ell = 0$, as they represent the net exchange of mass across the boundary of the system. A steady-state flux vector, $v \in \mathbb{R}^r$ such that $Sv = 0$, may be split into an internal reaction flux vector $z \in \mathbb{R}^n$ and an external reaction flux vector $w \in \mathbb{R}^k$, where $Nz + Bw = 0$.

We define a thermodynamically feasible flux as a net flux vector where each nonzero entry has a sign opposite change in chemical potential for the corresponding reaction, that is:

$$z_j > 0 \Rightarrow N_j^T y < 0 \quad (2)$$

$$z_j < 0 \Rightarrow N_j^T y > 0, \quad (3)$$

where one may interpret $N_j^T y$ as the thermodynamic driving force for reaction j and interpret $y \in \mathbb{R}^m$ as a vector proportional to the chemical potentials of each metabolite¹³. As there are fewer metabolites than reactions, any vector of chemical potentials imposes a set of sign constraints on net fluxes in order for the implications equations (2) and (3) hold. These equations are a relaxation of the strict thermodynamically feasible flux constraint $\text{sign}(z_j) = -\text{sign}(N_j^T y)$, because $z_j = 0 \Leftrightarrow N_j^T y = 0$.

Thermodynamically consistent model: we define a model to be thermodynamically consistent if each of its reactions admits a nonzero thermodynamically feasible flux. The largest set of reactions that are simultaneously thermodynamically feasible is defined by the cardinality optimization problem

$$\begin{aligned} & \max_{z, w, y} \|z\|_0 + \|w\|_0 \\ & \text{s.t. } Nz + Bw = 0, \\ & \quad \underline{z} \leq z \leq \bar{z}, \\ & \quad \underline{w} \leq w \leq \bar{w}, \\ & \quad z_j > 0 \Rightarrow N_j^T y < 0, \forall j \in 1 \dots n, \\ & \quad z_j < 0 \Rightarrow N_j^T y > 0, \forall j \in 1 \dots n, \end{aligned} \quad (4)$$

where $z \in \mathbb{R}^n$ is an internal reaction flux vector, $w \in \mathbb{R}^k$ is an external reaction flux vector and $y \in \mathbb{R}^m$ may be interpreted as a vector proportional to the chemical potentials of each metabolite¹³. Here $\underline{z}$ and $\bar{z}$ denote lower and upper bounds on internal reaction fluxes, whereas $\underline{w} \in \mathbb{R}^k$ and $\bar{w} \in \mathbb{R}^k$ denote lower and upper bounds on external reaction fluxes, respectively.

An algorithm for finding the largest thermodynamically feasible subset of a given model and its application to human metabolism is described in detail elsewhere¹¹. For completeness, here, we reiterate several key points that are relevant to the interpretation of this step. Inspired by

Desouki et al.²⁰, we observed that when the bounds on internal reactions constrain the directions but not magnitudes of net flux, that is $\underline{z} \in \{0, -\infty\}^n$ and $\bar{z} \in \{0, \infty\}^n$, a thermodynamically feasible flux can be computed by a single linear optimization problem that minimizes the absolute value of the internal net reaction flux.

As with flux consistency, typically, all reactions in a network do not simultaneously admit a nonzero thermodynamically feasible flux; therefore, an approximation to the largest thermodynamically flux consistent subset of reactions in a model can be found by amalgamating the reactions with at least one nonzero value in a set of thermodynamically feasible flux vectors. This set is obtained by a sequence of iterations that accumulates thermodynamically flux-consistent reactions by biasing net flux toward reactions that have not yet been identified as being thermodynamically flux consistent. This sequence of iterates is terminated when no further thermodynamically flux consistent reactions can be identified or a limit on the number of iterates is reached (by default, 30). The constraints in equation (4) define a nonconvex set; therefore, efficiently obtaining an approximation to the largest set of thermodynamically flux-consistent reactions is a pragmatic compromise. A convex set is a subset of a vector space (such as Euclidean space) that has a specific property: for any two points within the set, the line segment connecting these points lies entirely within the set. This convexity is crucial for computational efficiency and optimization in metabolic modeling, providing a robust foundation for refining models and achieving reliable solutions¹⁰.

The last two constraints in equation (4) are a relaxation of the strict thermodynamic constraint $\text{sign}(z_j) = -\text{sign}(N_j^T y)$ because $z_j = 0 \nleftrightarrow N_j^T y = 0$. However, they are sufficient to identify a thermodynamically flux-consistent subnetwork because any internal reaction with zero flux in all thermodynamically feasible flux vectors, which are each a solution to equation (4), is defined to be thermodynamically inconsistent.

The thermoKernel algorithm (Step 22)

We designed the `thermoKernel` algorithm to extract a context-specific model that is not only stoichiometrically and flux consistent but also thermodynamically flux consistent.

The core principle of `thermoKernel` is that it balances user-specified evidence of metabolite presence or absence and reaction activity or inactivity with thermodynamic feasibility constraints. The algorithm iteratively identifies feasible flux vectors that maximize the inclusion of desired features (for example, metabolites or reactions) while penalizing inclusion of undesired ones.

The key features of `thermoKernel` are as follows.

- It generates models consistent with omics and bibliomic data, weighted by user-specified importance
- It ensures that all included reactions admit at least one thermodynamically feasible flux
- It accommodates both qualitative (present, or absent) and quantitative (weighted) data
- It iteratively refines the model to guarantee that every metabolite is produced and consumed, and every retained reaction is active under at least one feasible flux vector

This algorithm represents a major advance over existing extraction methods by guaranteeing the thermodynamic consistency of the resulting models, thereby enabling more reliable predictions of cellular metabolic states.

Production and consumption of metabolites: `thermoKernel` assumes that a metabolite is present if it is either produced or consumed at a nonzero rate by a corresponding *net* flux vector and absent otherwise. Any stoichiometric matrix may be split into $N = R - F$, where $F_{i,j}$ and $R_{i,j}$ are the stoichiometric numbers of the *i*th molecule consumed and produced in the *j*th directed reaction, respectively. In terms of forward reaction flux, $p \in \mathbb{R}_{\geq 0}^n$, and reverse reaction flux, $q \in \mathbb{R}_{\geq 0}^n$, the rate of consumption of each metabolite is $Fp + Rq$ and the rate of production of each metabolite is $Fq + Rp$. Therefore, we observe that $s := (F + R)p + (F + R)q \geq 0$ represents the sum of the rate of production and consumption of each metabolite. If metabolite *i* is not present in a network, then $s_i = 0$.

Weighting of metabolites and reactions: with `thermoKernel`, you can generate a context-specific model by providing lists of present metabolites, absent metabolites, active reactions,

inactive reactions and weighted combinations thereof. You can generate a context-specific model only by providing a generic model and a list of present metabolites, a list of active reactions or a combination of both. With only a list of present metabolites, providing the generic model is thermodynamically flux consistent, a context-specific model is generated containing each of the metabolites required to be present. In this context, there is no constraint on which reactions must be active, except that at least one reaction must be active that results in a nonzero rate of production plus the consumption of each metabolite required to be present. With only a list of active reactions, providing the generic model is thermodynamically flux consistent, a context-specific model is generated containing each of the reactions required to be active. In this context, the list of metabolites that must be present is fully constrained by the list of reactions required to be active. When a combination of lists of present and absent metabolites as well as active and inactive reactions are provided, the resulting context-specific model is a tradeoff between these specifications, where, by default, equal weight is given to metabolite presence, metabolite absence, reaction activity and reaction inactivity.

Model extraction: `thermoKernel` identifies a context-specific yet thermodynamically consistent model from a model that is flux or also thermodynamically flux consistent. Each iteration of `thermoKernel` identifies a thermodynamically feasible flux vector that simultaneously quantitatively balances incentives for the presence of metabolites and the activity of reactions, with disincentives for the absence of metabolites and inactivity of reactions, utilizing the following optimization approach. This is represented by the following optimization problem:

$$\begin{aligned} \max_{z, w, p, q, s, r} \quad & \beta 1^T(p + q) + g^T \|z\|_0 + h^T \|s\|_0 + f^T \|r\|_0 \\ \text{s.t.} \quad & Nz + Bw = 0, \\ & z = p - q, \\ & (F + R)(p + q) = s, \\ & (F + R)z = r, \\ & \underline{z} \leq z \leq \bar{z}, \\ & \underline{w} \leq w \leq \bar{w}, \\ & 0 \leq p, \\ & 0 \leq q. \end{aligned} \tag{5}$$

Equation (5) contains the following variables:

- $z \in \mathbb{R}^n$ is an internal reaction flux vector
- $w \in \mathbb{R}^k$ is an external reaction flux vector
- $p \in \mathbb{R}_{\geq 0}^n$ is forward net reaction flux
- $q \in \mathbb{R}_{\geq 0}^n$ is reverse net reaction flux
- $s \in \mathbb{R}_{\geq 0}^m$ is the sum of the rate of production and consumption of each metabolite
- $r \in \mathbb{R}^m$ is an approximation to the sum of production and consumption of each metabolite due to net reaction flux

Equation (5) contains the following data matrices:

- $N \in \mathbb{N}^{m \times n}$ and $B \in \mathbb{N}^{m \times k}$ are internal (stoichiometrically consistent) and external (stoichiometrically inconsistent) stoichiometric matrices
- $F \in \mathbb{N}^{m \times n}$ and $R \in \mathbb{N}^{m \times n}$ are forward and reverse stoichiometric matrices¹¹, with $F = \max(-N, 0)$
- It contains $R = \max(N, 0)$ and $N = R - F$

Equation (5) contains the following data vectors:

- $\underline{z} \in \{0, -\infty\}^n$ and $\bar{z} \in \{0, \infty\}^n$ denote lower and upper bounds on internal reaction fluxes, respectively
- $\underline{w} \in \mathbb{R}^k$ and $\bar{w} \in \mathbb{R}^k$ denote lower and upper bounds on external reaction fluxes, respectively

In equation (5), the objective term $\beta 1^T(p + q)$ minimizes the absolute value of net internal reaction flux $z \in \mathbb{R}^n$, which optimizes toward a thermodynamically feasible flux vector²⁰. The scalar weight $\beta > 0$ is chosen to trade off between cardinality optimization of internal reaction fluxes, which explores the flux-consistent subset of the model, and minimizing the one norm

of the sum of forward and reverse reaction rates, which promotes thermodynamic feasibility. The parameter vector $g \in \mathbb{R}^n$ is a vector of weights on internal reaction rates. The term $g^T \|z\|_0$ denotes the weighted zero-norm of internal reaction fluxes, where the zero-norm for each reaction is weighted individually, that is:

$$g^T \|z\|_0 := \sum_{j=1}^n g_j \times \|z_j\|_0,$$

when $g_j < 0$, or $g_j > 0$, then net flux is incentivized or penalized, in reaction j , respectively.

In equation (5), the parameter vectors $h, f \in \mathbb{R}^m$ are vectors of weights on metabolites. When $h_j > 0$, the objective term $h_j^T \|s_i\|_0$ minimizes the zero-norm of the rate of production plus the consumption of metabolite i , promoting its elimination from a model. When $h_j < 0$, the objective term $h_j^T \|s_i\|_0$ maximizes the zero-norm of the rate of production and consumption of metabolite i . However, even if $s_i > 0$, then a metabolite may not be produced or consumed at a nonzero rate by a corresponding net flux vector because $p = q \neq 0$ is a solution that also implies $s > 0$. Therefore, we introduce the variable $r := (F + R)z$ as an approximation to the sum of production plus consumption due to net reaction flux, which is $(F + R)|z|$, but is nontrivial to implement. If $r_i := (F + R)_i z > 0$, then $(F + R)_i |z| > 0$ as desired. Therefore, when $f_i < 0$, the objective term $f_i^T \|r_i\|_0$ approximates the maximization of the weighted zero-norm of the sum of production plus consumption of metabolite i due to nonzero net reaction flux. In practice, we set $f_i = -\min(h_i, 0)$ so that optimization of the cardinality of s_i and r_i is concordant.

In equation (5), the internal bounds $\underline{z} \in \{0, -\infty\}^n$ and $\bar{z} \in \{0, \infty\}^n$ enforce reaction directionality without interfering with the minimization of the absolute value of internal net flux. That is, if $g = h = f = 0$ then equation (5) is guaranteed to return a thermodynamically feasible flux. However, we find that with $g = h = f \neq 0$, equation (5) still returns a thermodynamically feasible flux provided that the value of the strictly positive scalar parameter β is such that $\beta 1^T (p + q) \geq g^T \|z\|_0 + h^T \|s\|_0 + f^T \|r\|_0$ at the optimum of equation (5). In this way, equation (5) optimizes toward thermodynamically feasible flux and simultaneously quantitatively trades off incentives for presence of metabolites and activity of desired reactions, with penalties for the inclusion of undesired metabolites and reactions.

The following optional variants of equation (5) are implemented in `thermoKernel` depending on the data included in the preceding steps:

- A linear objective (Step 3) on net flux may be added to the objective in equation (5) to optimize a linear combination of net fluxes, as in flux balance analysis¹⁴. This will not interfere with thermodynamic feasibility if all the reactions being optimized are stoichiometrically inconsistent, that is, exchange reactions. However, it may be that the application of thermodynamic feasibility constraints reduces the maximum value of the part of the objective that optimizes net flux
- Coupling constraints³¹ (Step 15), $Cv \leq d$, may be added to equation (5), but depending on how they are formulated, such constraints may interfere with the objective term $\beta 1^T (p + q)$ that optimizes toward thermodynamically feasible flux. In that case, the coupling constraints may render part of the extracted model thermodynamically infeasible (for example, coupling constraints may be formulated that force net flux around a stoichiometrically balanced cycle)
- The box constraints net flux in equation (5) may be used to force nonzero flux in a particular direction of one or more internal reactions, that is $\underline{z}, \bar{z} \in \mathbb{R}^n$. Depending on the choice of bounds (Step 14), this may result in an extraction of a model where some reactions are thermodynamically inconsistent (for example, a combination of bounds that force net flux around a stoichiometrically balanced cycle)
- In `thermoKernel`, the implementation of equation (5) and its variants uses the default $\beta := 1 \times 10^{-4}$. By setting `model.beta`, you can set a larger value of β , which will increase the incentive to find a fully thermodynamically feasible flux vector in each iteration of equation (5) and decrease the incentive to search the steady-state solution space for a flux vector that results in certain reactions and metabolites to be active and present, respectively. However, if β is too large, relative to the magnitudes of the weights on metabolites and reactions, this impedes the ability of equation (5) to search the full set of thermodynamically feasible fluxes

- Similarly, there is a tradeoff between the reaction weights $g \in \mathbb{R}^n$ and the metabolite weights $h, f \in \mathbb{R}^m$, where we set $f_i = -\min(h_i, 0)$; therefore, their relative magnitude should reflect the relative importance of the corresponding input bibliomic, transcriptomic and proteomic data (Step 7)

Although it is not always necessary, the default `thermoKernel` implementation for each iteration of equation (5) should be followed by the linear optimization problem first introduced in previous work²⁰. This minimizes the absolute value of internal reaction flux, while (1) all exchange fluxes are kept constant, and (2) no internal flux is allowed to change direction or increase in size. This has two advantages. First, it is an efficient test to check that equation (5) returns a solution where every reaction corresponds to a thermodynamically feasible flux. Second, it can be used to remove that part of a flux vector that corresponds to flux around a stoichiometrically balanced cycle, and the output is a flux vector that is thermodynamically feasible, that is, each reaction flux satisfies equations (2) and (3). When `XomicsToModel` calls `thermoKernel`, the default behavior is for flux bounds, which interfere with thermodynamic feasibility, to be relaxed when testing for cyclic flux after equation (5). This prioritizes the thermodynamic consistency of models generated by `XomicsToModel` over user input constraints that attempt to force a contradiction.

Extracting a context-specific and thermodynamically consistent model: given a generic model, it is usually the case that more than one thermodynamically feasible flux vector is required to ensure the activity of a specified subset of its reactions, the presence of a specified subset of its metabolites or both. That is, a single thermodynamically feasible flux vector is usually unable to enforce the activity and presence of all the reactions and metabolites in a model. Therefore, a randomly greedy strategy is implemented, inspired by previous work⁴⁰, whereby it chooses a random vector from a uniform rectangular set, $d \in \{\mathbb{R}^n | \text{unif}(0, 1)\}$, then updates the weights on cardinality optimization of reactions in the next iterate $g(n+1)$ based on those of the previous iterate $g(n)$, with the heuristic:

$$g_j(n+1) := \begin{cases} g_j(0) \times (n+1) & g_j(n) < 0, z_j(n) = 0, d \geq 0.5, \\ 0 & g_j(n) < 0, z_j(n) = 0, d < 0.5, \\ 0 & g_j(n) < 0, z_j(n) \neq 0, \\ g_j(0) & g_j(n) > 0, \\ 0 & g_j(n) = 0. \end{cases} \quad (6)$$

This approach successively increases the incentive for activity of internal reactions that are desired to be active but have not yet been active in a previous iterate. The strategy is the same to incentivize the presence of metabolites. The iterations either conclude when all incentivized reactions and metabolites are active and present, respectively, or when a prespecified maximum number of iterations is reached. As reaction j is defined to be active when $z_j \neq 0$ (numerically implemented by requiring that $|z_j| \geq \varepsilon$ with $\varepsilon = 1 \times 10^{-4}$ in practice), every reaction added to a thermodynamically consistent model is a reaction that admits at least one nonzero thermodynamically feasible flux, and every metabolite is produced and consumed by at least one thermodynamically feasible flux. Therefore, while generating a single thermodynamically feasible flux vector containing nonzero and zero fluxes, the strict thermodynamically feasible flux constraint $\text{sign}(z_j) = -\text{sign}(N_j^T y)$ is relaxed; a thermodynamically consistent model is only generated from reactions that admit nonzero thermodynamically feasible fluxes. Therefore, the model generated is fully consistent with thermodynamics in its broadest sense. Each iteration of equation (5) is efficiently solved using a piecewise linear sequence of linear optimization problems, as described in detail elsewhere¹¹.

Limitations

This Protocol focuses on extracting a context-specific genome-scale model from an existing reconstruction by integrating omics data from a biological system. It does not cover the generation of a generic, genome-scale metabolic reconstruction from scratch^{2,4} or the analysis of existing models³⁴. Moreover, the integration of omics data presents several challenges

because they are generated using a variety of platforms, affecting storage and data formats substantially. The integration of omics data necessitates data in specific formats. Therefore, individual omics data must be preprocessed. Furthermore, experimental errors, such as data processing or measurement errors, can propagate through the extraction of a metabolic network that is not a faithful representation of the original biochemical system. In this situation, a significant loss of predictive capacity will be evident. Therefore, it is not recommended to extract a context-specific model and claim that it has high predictive accuracy without commensurate comparison with some independent experimental data.

We define a steady-state flux vector to be thermodynamically feasible if it is consistent with energy conservation and the second law of thermodynamics^{16,17}. They are thermodynamic constraints that apply to chemical and biochemical systems alike. In addition, biochemical systems are subject to many additional constraints that range from abiotic constraints⁴¹ to constraints imposed by evolutionary pressures toward optimization⁴². For example, there are biochemical thermodynamic and enzyme kinetic constraints on the range of metabolite concentrations required for life²⁴. Limits to the feasible range of metabolite concentrations lead to additional constraints on the direction⁴³ and magnitude⁴⁴ of reaction flux that are not applied within the XomicsToModel pipeline. The models generated by XomicsToModel are thermodynamically flux consistent; they can be integrated with additional constraints arising from biochemical thermodynamics and kinetics without a risk of generating a thermodynamically infeasible model. Several studies support this approach. For instance, one study⁴⁵ demonstrated the necessity of thermodynamic consistency in cyclic enzyme-catalyzed reaction networks, highlighting the strong constraints on kinetic parameters imposed by zero net flux cycles. Other groups⁴⁶ developed a formalism to combine kinetic and thermodynamic measurements into state-resolved models that are thermodynamically consistent, correcting potential inconsistencies in the models. These approaches ensure that the models are both accurate and feasible within the constraints of thermodynamic laws, supporting the assertion that XomicsToModel can handle additional biochemical constraints without issues.

Materials

Equipment

Input data

The COBRA Toolbox supports commonly used data formats for model description such as Systems Biology Markup Language and Excel sheets (.xlsx). MATLAB supports a wide range of text and spreadsheet formats that can be used to provide context-specific data for the XomicsToModel pipeline.

Required hardware

- A computer with at least 8 GB of RAM and any 64-bit Intel or AMD processor
▲ **CRITICAL** Depending on the size of the model, more processing power and more memory might be needed.
- A hard drive with at least 10 GB of space available

Required software

- An operating system that is MATLAB qualified (<https://mathworks.com/support/sysreq.html>)
▲ **CRITICAL** To make sure that the operating system is compatible with the MATLAB version, check the requirements at https://mathworks.com/support/sysreq/previous_releases.html.
- MATLAB (MathWorks, <https://mathworks.com/products/matlab.html>), version R2021+. Install MATLAB and its license by following the official installation instructions (<https://mathworks.com/help/install/ug/installmathworks-software.html>)
▲ **CRITICAL** Version R2021+ or above is recommended for running MATLAB live scripts (.mlx files).

Protocol

▲ **CRITICAL** No support is provided for versions older than R2021. MATLAB is released on a twice-yearly schedule, and the latest releases of MATLAB may not be compatible with the existing solver interfaces, necessitating an update of the MATLAB interface provided by the solver developers, an update of the COBRA Toolbox or both.

- The COBRA Toolbox version 3.4³⁴, or above. To install the COBRA Toolbox, follow the instructions on <https://github.com/opencobra/cobratoolbox>

▲ **CRITICAL** If an installation of COBRA Toolbox is already present, update the repository from MATLAB, via terminal, or Git Bash.

▲ **CRITICAL** Check that all of the system requirements in <https://opencobra.github.io/cobratoolbox/docs/requirements.html> are met.

Solvers

The table below provides an overview of the optimization solvers supported by the XomicsToModel pipeline.

Name	Version	Interface
GLPK	2.7+	glpk
GUROBI	7.0+	gurobi
ILOG CPLEX	10-12.10	ibm_cplex
MATLAB	R2014b+	matlab
MOSEK	8.0+	mosek

Equipment setup

COBRA Toolbox

Unless you are already a developer of COBRA Toolbox code, update the COBRA Toolbox from within MATLAB by running the following command.

```
>> updateCobraToolbox
```

Alternatively, update from the terminal (or Git Bash) by running the following from within the COBRA Toolbox directory.

```
$ cd cobratoolbox # change to the local cobratoolbox directory
$ git checkout master # switch to the master branch
$ git pull origin master # retrieve changes
```

If the update of the COBRA Toolbox fails or cannot be completed, clone the repository again.

▲ **CRITICAL** The COBRA Toolbox can be updated as described above only if no changes have been made to the code in one's local clone of the official online repository. To contribute any local edits to the COBRA Toolbox code into the official online repository, please follow the MATLAB.devTools guidelines³⁴ (<https://github.com/opencobra/MATLAB.devTools>).

To generate a context-specific model from a generic model (or reconstruction) using omics data, first launch the COBRA Toolbox³⁴ and specify the optimization solver to be used.

```
>> initCobraToolbox
>> changeCobraSolver('gurobi', 'LP');
>> changeCobraSolver('gurobi', 'QP');
```

Procedure

▲ **CRITICAL** The sequence of steps described below are in the same order as they are implemented by XomicsToModel. This is important because the overall approach to model generation is to first increase the dimensions and magnitude of the feasible set of steady-state

metabolic fluxes, then successively decrease its dimensions and magnitude, until a context-specific model and information about the procedure and result of the model generation process is returned to the user after the final step. The `XomicsToModel` pipeline can be run if at least one input variable and a generic model (in standard COBRA Toolbox format `COBRAModelFields.md`) is provided. The user-defined context-specific information (`specificData`) and technical parameters (`param`) are optional, and if not specified, default values for the required fields will be assigned.

```
>> [contextSpecificModel, modelGenerationReport] =  
XomicsToModel(genericModel, specificData, param)
```

The name and types of the field for the `specificData` and `param` variables must be identical to how they are described in each step of the `XomicsToModel` pipeline below so that the `XomicsToModel` pipeline can recognize them. The supplementary `preprocessingOmicsModel` function can be used to generate the `specificData` variable, as described in the Supplementary Information.

Data preparation

▲ **CRITICAL** In this first step, `XomicsToModel` checks the context-specific data in `specificData`, checks the technical parameters in `param` and checks the generic model to test if they are compatible with the naming conventions and formats expected by the rest of the `XomicsToModel` pipeline. `XomicsToModel` identifies missing `specificData` and `param` fields, and default values are assigned for optional inputs.

1. To aid interpretation of the performance of the pipeline, specify a filename for the diary in `param.diaryFilename`. To enhance the interpretation of the `XomicsToModel` pipeline performance, a verbose level can be implemented in `param.printLevel`. To control the amount of information displayed for each step, adjust the verbose level parameter from 0 to 3. A verbose level of 0 provides minimal output, focusing only on essential information (but not completely silent), whereas a level of 1 includes additional warnings and status updates. At level 2, more detailed process information is displayed, including intermediate steps, and at level 3, the most comprehensive output is provided, showing detailed logs and debugging information.

Within this diary file, many intermediate results of the pipeline are printed that are essential for debugging. If a certain reaction is missing that was expected within the generated model, search for it in the diary file. Specify `param.debug = 1` to print out intermediate results after key steps of the `XomicsToModel` pipeline. Later steps in the procedure will mention that results are printed; to find these results, open the diary.

These intermediate results can be reloaded so that subsequent individual steps of the `XomicsToModel` pipeline may be run to debug unexpected results. Alternatively, these intermediate results can be analyzed as a batch to monitor the model generation progress, as described in the Supplementary Information.

Usable variables for Step 1

`param.debug`: logical, true if the pipeline should save its progress for debugging (default: false).

`param.diaryFilename`: the name (and location) of a diary file with the printed pipeline output (default: 0).

`param.printLevel`: level of verbose that should be printed (default: 0).

Generic model checks

▲ **CRITICAL** In this step, the generic model is checked for compatibility with `XomicsToModel`. It is checked for fields that may cause inconsistency when reactions are added or removed by `XomicsToModel`. These checks depend on the model extraction algorithm used by `XomicsToModel`.

2. Select your preferred model extraction algorithm by setting `param.tissueSpecificSolver`. The default is 'thermoKernel', which will ensure

the model extracted admits a thermodynamically feasible flux for each reaction.

Alternatively, set `param.tissueSpecificSolver = 'fastcore'`.

Regardless of the model extraction algorithm selected, every generic model must satisfy the steady state constraints $Sv = 0$, reaction rate bounds $l \leq v \leq u$ and optionally coupling constraints $Cv \leq d$, implying that the system of inequalities is feasible; if the model is infeasible in this step, an error will be generated.

Review the log; if `param.printLevel` is greater than zero, the information in `specificData` and `param` will be printed.

◆ TROUBLESHOOTING

Usable variables for Step 2

`param.printLevel`: level of verbose that should be printed (default: 0).

`param.tissueSpecificSolver`: the name of the tissue-specific solver to be used to extract the context-specific model (possible options: 'thermoKernel' and 'fastcore'; default: 'thermoKernel').

Set objective function (optional)

▲ **CRITICAL** `XomicsToModel` has the capability to set an objective function as used in flux balance analysis¹⁴, where $c \in \mathbb{R}^n$

represents the biologically inspired linear objective to find the optimal flux vector.

3. (Optional) Unless you are using default settings, set the linear objective $\varphi(v) := c^T v$ in `param.setObjective`. The objective function in `param.setObjective` represents the reaction whose flux is to be maximized (such as adenosine triphosphate (ATP) or biomass production) or minimized (such as energy consumption). However, alternative biological objectives can also be used. For example, these may include coupled reactions or mathematical objectives such as the minimization of zero, one or two norms; entropy maximization; or a combination of biological and mathematical objectives. Examples of such combinations include the minimization of norms combined with reaction expression data based on gene expression or the minimization of bonds broken and formed based on atom mapping data at a genome scale. These approaches are illustrated in Preciat et al.⁹.

Usable variable for Step 3

`param.setObjective`: reaction identifier indicating the reaction to maximize (default: none).

Add missing reactions (optional)

▲ **CRITICAL** If a reaction is already present in the draft model, it will be replaced by the new reaction list.

▲ **CRITICAL** By associating a new gene identifier with a reaction, a new gene will be added to the draft model.

4. (Optional) If reactions are missing from the generic model, specify the missing reactions to add in `specificData.rxns2add`. These reactions are then automatically added to the draft model by `XomicsToModel` using the function [addReaction](#). When specifying missing reactions, multiple data can be included but only the reaction identifier and the reaction formula are mandatory.

Add missing metabolic reactions on `specificData.rxns2add`. The default bounds are set on the basis of the arrows used in the reaction formula to separate the substrates and products as shown in the table `specificData.rxns2add` in the 'Usable variables for Step 4' section (below). The available arrows are:

- Forward ($\rightarrow$): $l_j = 0$; $u_j = \text{param.TolMaxBoundary}$
- Reverse ($\leftarrow$): $l_j = \text{param.TolMinBoundary}$; $u_j = 0$
- Reversible ($\leftrightarrow$): $l_j = \text{param.TolMinBoundary}$; $u_j = \text{param.TolMaxBoundary}$

If not otherwise specified, the default values for missing reaction names, metabolic pathways and gene rules are custom reaction, miscellaneous and an empty cell, respectively.

Adding reactions to a model may render it stoichiometrically inconsistent. Given a generic reconstruction or model, `XomicsToModel` generates a generic stoichiometrically consistent model using the function `findStoichConsistentSubset` described elsewhere¹¹. In brief, this function approximately solves the problem to extract the largest subset of a reconstruction or model that is stoichiometrically consistent³⁴. A set of reactions is stoichiometrically consistent if every reaction in that set is mass balanced⁴⁷. The stoichiometric matrix may be split $S := [N, B]$, where the matrix N represents stoichiometrically consistent internal reactions, whereas the matrix B represents stoichiometrically inconsistent external reactions.

Review the log; if `param.printLevel` is greater than zero, a table with a summary of the draft model's stoichiometric consistency will be printed (Table 3).

The inconsistent metabolites and reactions are removed; if this renders a model infeasible with respect to steady-state flux, an error is generated by `XomicsToModel`. In addition, in this step, the gene–protein–reaction rules are updated for the newly added reactions, and the reaction–gene–matrix is regenerated. In addition, different Boolean vectors indicating stoichiometrically consistent and inconsistent metabolites and reactions are added to the draft model.

▲ **CRITICAL STEP** Extracting the largest stoichiometrically consistent subset is combinatorially complex. At the genome-scale, only approximation to the actual solution is achievable in practice. The quality of the approximation is improved if `findStoichConsistentSubset` is warm-started with reactions specified as being stoichiometrically inconsistent (for example, external reactions specified as false in the corresponding entry in `model.SIntRxnBool`). If `model.SIntRxnBool` is not provided, `XomicsToModel` calls the function `findSExRxnInd` to heuristically distinguish internal and external reactions on the basis of reaction-naming conventions (for example, 'EX_' prefix for an exchange reaction) and reaction stoichiometry (any reaction with only one nonzero stoichiometric coefficient in the corresponding row is an external reaction).

◆ TROUBLESHOOTING

Usable variables for Step 4

`model.SIntRxnBool`: logical vector, true if a reaction is heuristically considered an internal reaction, false for an external reactions (default: false).

`specificData.rxns2add`: table containing the identifier of the reaction to be added, its name, the reaction formula, the metabolic pathway to which it belongs, the gene rules to which the reaction is subject and the references (default: empty).

rxns	rxnNames	rxnFormulas	subSystems	rxnGrRules	rxnReferences
newRxn1	Oxidation	$A + O \rightarrow AO$	Glycolysis	gene1 or gene2	PMID: ****
newRxn2	Reduction	$AH + B \leftrightarrow A + BH$	Glycolysis	gene1	PMID: ****

`param.TolMaxBoundary`: the reaction-flux upper-bound maximum value (default: 1,000).

`param.TolMinBoundary`: the reaction-flux lower-bound minimum value (default: -1,000).

Identify core metabolites and reactions (optional)

▲ **CRITICAL** The `XomicsToModel` pipeline uses several user-specified fields from the `specificData` variable to identify the set of metabolites and reactions that must be active in the draft model. These optional fields include the reaction specified as the objective function, the metabolic reactions to add, the metabolites and reactions identified as active based on bibliomic data, the reactions whose bounds are restricted based on bibliomic data, the coupled reactions, the exchange reactions for metabolites present in the growth media and the exchange reactions for metabolites shown to be secreted or taken up based on exometabolomic data.

- (Optional) Add reactions to the active reaction list, which will be appended or shortened if necessary to keep the model stoichiometrically consistent and flux consistent (and optionally thermodynamically consistent). To configure these fields, perform the following actions and refer to the usable variables table for syntax:

- Activate specific reactions in `specificData.activeReactions` based on bibliomic data

- Specify a linear objective function in `param.setObjective` (Step 3)
- Generate a table of reactions to be added in `specificData.rxns2add` (Step 4)
- List the metabolites known to be absent in `specificData.absentMetabolites` based on the bibliomic data
- List of metabolites known to be present in `specificData.presentMetabolites` based on bibliomic data
- Include a list of metabolites present in the cell culture in `specificData.mediaData` if the omics data originated from one.
- Utilize qualitative or semiquantitative exometabolomic data in `specificData.exomet`
- Formulate a list of reactions for constraint in `specificData.rxns2constrain` based on bibliomic data
- Establish a set of coupled reactions in `specificData.coupledRxns`

Usable variables for Step 5

`param.setObjective`: cell string indicating the linear objective function to optimize (default: none).

`specificData.activeReactions`: list of reactions know to be active based on bibliomic data (default: empty).

`specificData.coupledRxns`: table containing information about the coupled reactions. This includes the identifier for the coupling constraint (`couplingConstraintID`), the list of coupled reactions (`coupledRxnsList`), the coefficients of those reactions (`c`, given in the same order), the right hand side of the constraint (`d`), the constraint sense or the directionality of the constraint (`dsense`) and the reference (default: empty).

couplingConstraintID	coupledRxnsList	c	d	d _{sense}	Reference
cRxn1	crxn1, crxn2	1 1	2.0250	G	PMID: ****
cRxn2	crxn1, crxn3	1 -1	0.2265	G	PMID: ****

`specificData.exomet`: table with measured exchange fluxes, for example, obtained from an exometabolomic analysis of fresh and spent media. It includes the reaction identifier, the reaction name, the measured mean flux, standard deviation of the measured flux, the flux units and the platform used to measure it (default: empty).

rxns	rxnNames	Mean	s.d	Units	Platform
EX_mMet1	Exchange of mMet1	100	8×10^{-4}	$\mu\text{mol/gDW/h}$	Liquid chromatography–mass spectrometry (LC–MS)
EX_mMet2	Exchange of mMet2	4×10^{-3}	3×10^{-4}	$\mu\text{mol/gDW/h}$	Gas chromatography–mass spectrometry (GC–MS)

`specificData.mediaData`: table containing the fresh media concentrations. It contains the reaction identifier; the maximum uptake ($\mu\text{mol/gDW/h}$) assuming it is equal to the concentration of the metabolite in fresh medium divided by the length of time before medium exchange; and the medium concentration ($\mu\text{mol/l}$; default: empty).

rxns	mediumMaxUptake	mediumConcentrations
EX_met1	–220	1,345
EX_met2	–14,916	191,021

`specificData.rxns2add`: table containing the identifier of the reaction to be added, its name, the reaction formula, the metabolic pathway to which it belongs, the gene rules to which the reaction is subject and the references (default: empty).

rxns	rxnNames	rxnFormulas	subSystems	rxnGrRules	rxnReferences
newRxn1	Oxidation	$A + O \rightarrow AO$	Glycolysis	gene1 or gene2	PMID: ****
newRxn2	Reduction	$AH + B \leftrightarrow A + BH$	Glycolysis	gene1	PMID: ****

`specificData.rxns2constrain`: table containing the reaction identifier, the updated lower bound, the updated upper bound, a constraint description and any notes such as references or special cases (default: empty).

rxns	Lower bound	Upper bound	Constraint description	Notes
rxn1	0		Close uptake	PMID: ****
rxn2	0.956	1,000	Turnover constraints	PMID: ****

Set limit bounds (optional)

- (Optional) Considering the constraints $l \leq v \leq u$ shown in Equation (1), where l is the lower bound and u the upper bound, set `param.TolMinBoundary` so that any minimum lower bound is assigned to the minimum value of the lower bounds. Similarly, set `param.TolMaxBoundary` so that any maximum upper bound is assigned to the maximum value of the upper bounds. After this adjustment, $\min(l) = \text{param.TolMinBoundary}$ and $\max(u) = \text{param.TolMaxBoundary}$.

To modify the lower and upper boundaries of the reaction fluxes, set:

- `param.TolMinBoundary = min(l)`
- `param.TolMaxBoundary = max(u)`

`param.TolMinBoundary` represents the minimum value a reaction can have, ensuring that no reaction in the network will have a flux lower than this threshold. Similarly, `param.TolMaxBoundary` sets the maximum allowable flux value for reactions, preventing fluxes exceeding this threshold in the entire network. The model's feasibility is tested with the new bounds.

Review the log; if `param.printLevel` is greater than zero, the modified bounds will be printed.

Usable variables for Step 6

`param.TolMaxBoundary`: the reaction-flux upper-bound maximum value (default: 1,000).

`param.TolMinBoundary`: the reaction-flux lower-bound minimum value (default: -1,000).

▲ **CRITICAL STEP** There should be no lower bound that exceeds an upper bound.

Identify active genes (optional)

▲ **CRITICAL** A gene must be present in the draft model to be considered active.

- Specify a list of active genes by listing their NCBI Entrez gene identifiers, which are unique integers for genes or other loci, or by using the gene identifier employed in the generic model. Compose the active genes list on the basis of data from multiple sources, such as bibliomics, proteomics and/or transcriptomics.
 - Specify which genes should be added using the gene–protein–reaction rules in `specificData.rxns2add.rxnGrRules` (Step 4)
 - Bibliomics: classify genes as active in `specificData.activeGenes` on the basis of a curation of published literature using your preferred specifications
 - Proteomics: classify genes as active on the basis of their (normalized) proteomic peak areas from the list of detected proteins (`specificData.proteomics`). If $\text{peakArea} \geq \text{param.thresholdP}$, the gene is considered active
 - Transcriptomics: add on `specificData.transcriptomics` genes that, on the basis of the list of detected transcripts, are active based on their expression levels (FPKM or raw counts). If $\log(\text{expressionLevel}) \geq \text{param.thresholdT}$, the gene is considered active

Set `param.inactiveGenesTranscriptomics = true` to set it so that the genes that fall below the specified threshold are added to the list `specificData.inactiveGenes`.

In the event of a discrepancy, such as an inactive gene according to manual literature curation but active according to proteomic or transcriptomic data:

- Set `param.curationOverOmics = true` to give preference to bibliomic data
- Set `param.curationOverOmics = false` to give preference to omics data

By default, bibliomic data takes precedence over other omics data, that is, `param.curationOverOmics = false`. The genes not found in the draft model are removed from the list of active genes.

Furthermore, two fields are added to the model to indicate the gene expression value of the genes present in the draft model `model.geneExpVal` and the expression of the corresponding reactions (`model.expressionRxns`) based on the function `mapExpressionToReactions`. The logarithm of the levels of gene expression in `specificData.transcriptomicData` are used to link gene expression levels to associated reactions using the function `mapExpressionToReactions`. In short, the resulting `model.expressionRxns` field is based on the gene relationship described in the gene-protein-reaction rules: an AND in the gene-protein-reaction rules is assigned $\min(l)$, and an OR is assigned a $\max(u)$. For genes that are not found in the transcriptomics data, no expression value is added (NaN is the default).

Although the default threshold for `specificData.proteomics` and `specificData.transcriptomicData` data is 0, a common practice in the field is to set it to $\log(2)$, corresponding to a fold change of 2. This threshold indicates that the expression level of a gene or protein is at least doubled compared with the baseline, highlighting a biologically meaningful increase and reducing noise from low-expressed genes or proteins⁴⁸.

Review the log; if `param.printLevel` is greater than zero, a histogram of reaction expression is printed.

Proteomics data are usually annotated with [UniProtKB](#). To integrate proteomics data with a generic model, it is essential to convert UniProtKB identifiers to gene identifiers using bioinformatics tools such as DAVID or UniProt's mapping services. This conversion allows seamless translation and facilitates comprehensive analyses across diverse biological databases.

Usable variables for Step 7

`param.curationOverOmics`: logical, indicates whether curated data should take priority over omics data (default: false).

`param.inactiveGenesTranscriptomics`: logical, indicates if inactive genes in the transcriptomic analysis should be added to the list of inactive genes (default: true).

`param.thresholdP`: the proteomic cutoff threshold for determining whether or not a gene is active (default: 0).

`param.thresholdT`: the transcriptomic cutoff threshold for determining whether or not a gene is active (default: 0).

`specificData.activeGenes`: list of Entrez identifiers of genes that are known to be active based on the bibliomic data (default: empty).

`specificData.inactiveGenes`: list of Entrez identifiers of genes known to be inactive based on the bibliomics data (default: empty).

`specificData.proteomics`: table with a column with Entrez identifiers and a column for the corresponding protein levels (default: empty).

Gene	expVal (expression value)
entrezID1	24
entrezID2	78

`specificData.transcriptomics`: table with a column with Entrez identifiers and a column for the corresponding transcriptomics expression value in FPKM or raw counts (default: empty).

Gene	expVal
entrezID1	220
entrezID2	14,916

- (Optional) `XomicsToModel` can close the lower and upper bounds for ion exchange reactions, preventing the net exchange of ions with the environment. To close the ions

in the model, set `param.closeIons = true`. By doing so, it is possible to defer the modeling of ion exchange (for example, to represent action potentials in neurons), as they represent a dynamic trajectory of ion concentrations, while prioritizing metabolic activities. Enabling ion exchange could be done where there is experimental evidence for net secretion or uptake of ions. The rate of uptake or secretion should then be specified for each ion in Step 14.

Usable variable for Step 8

`param.closeIons`: logical, true if ion exchange reactions are to be closed (default: false).

▲ **CRITICAL** If the parameter `param.closeUptakes` is set to true, ensure that uptakes are included; otherwise, it may lead to inconsistent models owing to the absence of uptake reactions.

9. (Optional) To close the flux of exchange reactions, set `param.closeUptakes = true`. When this action is taken, the lower bound of the exchange reactions in the draft model are closed (set to 0) except for those reactions subject to `specificData.rxns2constrain`. This will allow the draft model to only take up the metabolites specified by the bibliomics, growth media or metabolomics data.

Set `param.closeUptakes = true` for biological systems derived from defined cell cultures.

Review the log; if `param.printLevel` is greater than zero, the new constraints will be printed along with statistics such as the size of the stoichiometric matrix, the number of exchange reactions closed, the number of exchange reactions in active reactions and the number of exchange reactions in reactions to constrain based on manual literature curation.

Usable variables for Step 9

`param.closeUptakes`: logical, decide whether or not all of the uptakes in the draft model will be closed (default: false).

`specificData.rxns2constrain`: table containing the reaction identifier, the updated lower bound, the updated upper bound, a constraint description and any notes such as references or special cases (default: empty).

rxns	Lower bound	Upper bound	Constraint description	Notes
rxn1	0		Close uptake	PMID: ****
rxn2	0.956	1,000	Turnover constraints	PMID: ****

Close sink and demand reactions (optional)

10. (Optional) To set the demand and sink reactions to be closed, edit the parameter `param.nonCoreSinkDemands = true`. The demand and sink reactions specified in `specificData.activeReactions` will be exempt from closure. This distinction is vital, as closed reactions are not considered by the extraction algorithm for inclusion in the final model.

The solution of a feasible model must satisfy the constraints $Sv = 0$ and $l \leq v \leq u$ and optionally coupling constraints $Cv \leq d$. The newly added constraints may not all be consistent with a steady state flux at the same time, implying that the system of inequalities (and therefore the model) is infeasible, which could be caused by an incorrectly specified reaction bound. To perturb the constraints on an infeasible model to make it feasible, we use [relaxedFBA](#), a constraint relaxation function that uses cardinality optimization to find a minimal number of constraints to relax to make a model admit a steady-state flux vector (that is, a feasible flux balance analysis problem). Cardinality optimization is implemented by approximating it with an efficient continuous optimization algorithm¹¹.

Set the relaxation options in `param.relaxOptions` (for example, to only relax bound constraints or steady state constraints) for all or defined sets of reactions or metabolites.

`XomicsToModel` saves the old bounds from spent media metabolites in two new fields in the draft model `model.lbpreSinkDemandOff` and `model.ubpreSinkDemandOff`, along with fields identifying the reactions that were relaxed

([relaxedFBA](#)). `relaxedFBA` is a flexible function for relaxing constraints in `XomicsToModel`. The user may control it by setting the parameters in `param.relaxOptions` (for example, to relax steady state, flux bounds or subsets and combinations thereof). The possible options are described in detail in step 23 of COBRA Toolbox v3.0 protocol³⁴.

Review the log; if `param.printLevel` is greater than zero, the number and names of closed demand or sink reactions are printed.

Usable variables for Step 10

`param.nonCoreSinkDemands`: the type of sink or demand reaction to close is indicated by a string (possible options: 'closeReversible', 'closeForward', 'closeAll' and 'closeNone'; default: 'closeNone').

`param.relaxOptions`: a structure array with the relaxation options (default: `param.relaxOptions.steadyStateRelax = 0`).

`specificData.activeReactions`: list of reactions known to be active based on bibliomic data (default: empty).

Set metabolic constraints (optional)

▲ **CRITICAL** Extracellular metabolic constraints can be obtained from two different sources: fresh cell culture media metabolite concentrations (as specified by the manufacturer or as specified by combining known quantities of metabolically specified media), quantitative extracellular metabolomic measurements of fresh and spent media (assuming a metabolomic steady state) or both. The difference between spent and fresh media concentrations for extracellular metabolites is assumed to equal the rate of exchange of the metabolite between the cells and cell culture media (discussed further below). Based on your overall objective, constraints from fresh cell culture media metabolite concentrations or extracellular metabolomic measurements can either be integrated before or after the extraction algorithm step.

11. (Optional) To integrate constraints, set the `param.`

`growthMediaBeforeModelExtraction` and `param.`

`metabolomisBeforeModelExtraction` parameters respectively. Consider the following important factors when specifying the set of reactions to extract.

- The default is `param.growthMediaBeforeModelExtraction = true` and `param.metabolomisBeforeModelExtraction = true` as it helps to ensure that the model that enters (and leaves) the extraction step is consistent with metabolomic data
- Setting `param.growthMediaBeforeModelExtraction = false` and `param.metabolomisBeforeModelExtraction = false` usually results in a less constrained model, which can be useful if one wishes to change the exometabolomic constraints (for example, to model perturbations). However, care must be taken to ensure that reactions that exchange metabolites corresponding to exometabolomic measurements are still present in the extracted model

If the constraints are added before the extraction, the draft model solution space will be smaller. If it is done after the extraction, the context-specific model is extracted from a larger solution space with only the reaction directionality, steady-state and perhaps coupling constraints, given a particular network input into the model extraction step.

Usable variables for Step 11

`param.growthMediaBeforeModelExtraction`: logical, should the cell culture data be added before (true) or after (false) the model extraction (default: true).

`param.metabolomisBeforeModelExtraction`: logical, should the metabolomic data be added before (true) or after (false) the model extraction (default: true).

Cell culture data (optional)

▲ **CRITICAL** The uptake rates for exchange reactions corresponding to metabolites present in the media are adjusted by `XomicsToModel` on the basis of their concentration by modifying the

lower bounds on uptake rates, assuming the maximum uptake rate completely removes each metabolite from the medium between media replacements.

12. (Optional) Use the supplementary function `preprocessingOmicsModel`, described in the Supplementary Information, to calculate the maximum medium uptake rates on the basis of the fresh media metabolite concentrations. Set the field `mediumConcentrations` in `specificData.mediaData` to incorporate either the measured concentrations of metabolites in fresh media or those obtained from specifications of the media components. This will include the fresh media metabolite concentrations in subsequent calculations.

Usable variable for Step 12

`specificData.mediaData`: table containing the fresh media concentrations. Contains the reaction identifier, the maximum uptake ($\mu\text{mol/gDW/h}$) assuming it is equal to the concentration of the metabolite in fresh medium divided by the length of time before medium exchange and the medium concentration ($\mu\text{mol/l}$; Default: empty).

rxns	mediumMaxUptake	mediumConcentrations
EX_met1	-220	1,345
EX_met2	-14,916	191,021

Metabolomics (optional)

▲ **CRITICAL** Metabolomic exchange reaction constraints can be derived from quantitative metabolomic experiments. In the context of `XomicsToModel`, it is assumed that these measurements originate from the same biological sample at two different time-points. Assess the quality of the data carefully before integrating it with the model by following the general community standards (measurement accuracy, precision, blank effect, linearity of the calibration line, relative errors of calibration line fit and so on)⁴⁹.

▲ **CRITICAL** The measured metabolites should be mapped onto the model namespace where metabolite identifiers and/or exchange reaction identifiers for each metabolite should be found (ideally by cross-mapping standard metabolite identifiers such as InChI, ChEBI, HMDB and so on) and included in the final metabolomics input data.

13. (Optional) Specify the rates of consumption or production on the basis of exometabolomic data in `specificData.exoMet`. Outlier samples should be identified (for example, based on the median and quartiles of detected metabolite levels) and excluded. The metabolomic data are transformed into metabolic rates and normalized for the biomass content, for example, per gram of dry weight (gDW) with:

$$\nu_{\text{met}} := \frac{(\bar{x}_1 - \bar{x}_0) \times f}{\tau \times p}$$

where ν_{met} is the rate of exchange of a given metabolite, $\bar{x}_0$ is the mean initial metabolite concentration (for example, in $\mu\text{mol/L}$), $\bar{x}_1$ is the mean metabolite concentration at the second time point ($\mu\text{mol/L}$), τ is the time interval between the two measurements (in hours), p is the protein concentration in the sample (g/L) and f is the protein fraction of the total dry weight (g protein/gDW). It should be noted that the biomass content can also be estimated on the basis of the cell count and an estimate of the dry weight mass of a cell or by direct measurements of dry weight mass (tissue). To propagate the error associated with exchange flux measurement, it is possible to propagate the standard deviation associated with each measured value using:

$$\sigma_{\Delta x} := \sqrt{\sigma(\bar{x}_1)^2 + \sigma(\bar{x}_0)^2}, \sigma_{\nu} := \frac{|\nu_{\text{met}}| \cdot p \cdot f \cdot \sqrt{\left(\frac{\sigma(\Delta x)}{\Delta x}\right)^2 + \left(\frac{\sigma(p)}{p}\right)^2}}{\tau} \quad (7)$$

where $\sigma(\Delta x)$ is the standard deviation in the estimation of the change in metabolite concentration between initial ($\bar{x}_0$) and final measurement ($\bar{x}_1$) given the measurement uncertainties, $\sigma(\bar{x}_1)$ and $\sigma(\bar{x}_2)$, respectively, and σ_{ν} is the standard deviation in the exchange reaction rate, taking into account the standard deviation in metabolite concentration change estimation $\sigma(\Delta x)$ and standard deviation in protein estimation $\sigma(p)$. As the standard deviation in protein estimation is

the same for each metabolite, when fitting a model that is linear in the rate of each reaction, one can consider this to be zero. The metabolomic constraints are set by fitting the bounds of the draft model to the exometabolomic reaction rates while allowing the relaxation of net flux bounds based on the quadratic optimization problem presented in previous work⁹:

$$\begin{aligned} \min_{v, p, q \in \mathbb{R}^n} & (v_{\text{exp}} - v)^T \text{diag}(w_{\text{exp}})(v_{\text{exp}} - v) + p^T \text{diag}(w_l)p + q^T \text{diag}(w_u)q \\ \text{s.t. } & S = v0, \\ & Cv \leq d, \\ & l - p \leq v \leq u + q, \\ & 0 \leq p, \\ & 0 \leq q, \end{aligned} \quad (8)$$

where $p \in \mathbb{R}_{\geq 0}^n$ and $q \in \mathbb{R}_{\geq 0}^n$ are non-negative variables that allow the lower and upper bound constraints to be relaxed, l and u , respectively. The matrix–vector inequality $Cv \leq d$ represents coupling constraints, where the data $C \in \mathbb{R}^{s \times n}$ and $d \in \mathbb{R}^s$ enforce s coupling constraints between net reaction fluxes (for example, to represent a constraint on cumulative turnover of a metabolite by a set of degradation reactions). This problem always returns a steady-state flux $v \in \mathbb{R}^n$ and enables different information to be input as parameters, including the penalization of deviation from experimental fluxes ($w_{\text{exp}} \in \mathbb{R}_{\geq 0}^n$), penalizing relaxation of lower bounds ($w_l \in \mathbb{R}_{\geq 0}^n$) and penalizing relaxation of upper bounds ($w_u \in \mathbb{R}_{\geq 0}^n$).

Define the weights for penalizing the experimental fluxes in `param.metabolomicWeights`. By default, these weights are derived from the ‘mean’ of the experimental reaction flux. The weights could be derived from the standard deviation (`sd` in (2), ‘SD’) or relative standard deviation (‘RSD’) of the experimental reaction flux. For example, `param.metabolomicWeights = ‘SD’` sets the penalty on deviation from experimental measurement to be the inverse of one plus the variance:

$$v_{\text{exp}} := \frac{1}{(1 + \sigma_{\text{exp}}^2)} \quad (9)$$

where $\sigma_{\text{exp}} \in \mathbb{R}_{\geq 0}^n$ is the standard deviation in measured fluxes. This approach increases the penalty on deviation from an experimentally measured mean flux where the variance is lower.

To avoid numerical errors in model analysis, a default flux value is used if the experimental data exceed the specified boundary limits: the minimum reaction rate (`param.TolMinBoundary`), the maximum reaction rate (`param.TolMaxBoundary`) as well as the absolute minimum flux allowed, as defined in `param.boundPrecisionLimit`, that is, if any $|l_j|$ or $|u_j| < \text{param.boundPrecisionLimit}$, then it is considered as zero.

Review the log; if `param.printLevel` is greater than zero, the new metabolomically derived bounds on reaction fluxes are printed.

Finally, the feasibility of the draft model is tested, and if it is not feasible, for example, if there are inconsistencies or errors in the exometabolomic data, `XomicsToModel` will relax some reaction bounds using [relaxedFBA](#).

Usable variables for Step 13

`param.boundPrecisionLimit`: precision of flux estimate; if the absolute value of the lower bound or the upper bound is lower than the *boundPrecisionLimit* but higher than 0, the value will be set to the *boundPrecisionLimit* (default: primal LP feasibility tolerance).

```
>> feasTol = getCobraSolverParams('LP', 'feasTol');
>> param.boundPrecisionLimit = feasTol*10;
```

`param.metabolomicWeights`: string indicating the type of weights to be applied for metabolomics fitting (possible options: ‘SD’, ‘mean’ and ‘RSD’; default: ‘mean’).

`param.TolMaxBoundary`: the reaction-flux upper-bound maximum value (default: 1,000).
`param.TolMinBoundary`: the reaction-flux lower-bound minimum value (default: -1,000).
`specificData.essentialAA`: list exchange reactions of essential amino acids (default: empty). It must never be secreted, even in a relaxed flux balance analysis model.
`specificData.exomet`: table with measured exchange fluxes, for example, obtained from an exometabolomic analysis of fresh and spent media. It includes the reaction identifier, the reaction name, the measured mean flux, standard deviation of the measured flux, the flux units and the platform used to measure it (default: empty).

rxns	rxnNames	Mean	s.d	Units	Platform
EX_mMet1	Exchange of mMet1	100	8×10^{-4}	$\mu\text{mol/gDW/h}$	LC-MS
EX_mMet2	Exchange of mMet2	4×10^{-3}	3×10^{-4}	$\mu\text{mol/gDW/h}$	GC-MS

Add custom constraints (optional)

▲ **CRITICAL** A custom constraint enables the specification of precise bounds on reactions, derived from bibliomic data or the enforcement of essential reaction activity within the network.

14. (Optional) Specify bibliomic constraints in the struct array `specificData`.

`rxns2constrain` described in the 'Usable variables for Step 14' section to add both internal and external constraints on reactions. This information will be used to change the lower and upper bounds of the reactions in the draft model. Demand reactions will be ignored to ensure thermodynamic consistency if 'thermoKernel' is chosen as the tissue-specific solver. This means that constraints on demand reactions are not considered, which helps maintain the thermodynamic consistency of the model generation process. No change is made if the bound is empty. If the bounds in the table exceed the established limit bounds, they will be adjusted to what the user or default data specifies. Finally, the feasibility of the draft model is tested, and, if it is not feasible, `XomicsToModel` will relax some reaction bounds using [relaxedFBA](#).

Review the log; if `param.printLevel` is greater than zero, the custom constraints with nondefault bounds will be printed. If any discrepancies or issues are identified, refer to the 'Troubleshooting' section to resolve common constraint errors (for example, bounds exceeding limits).

◆ TROUBLESHOOTING

Usable variables for Step 14

`param.boundPrecisionLimit`: precision of flux estimate, if the absolute value of the lower bound or the upper bound is lower than the *boundPrecisionLimit* but higher than 0, the value will be set to the *boundPrecisionLimit* (default: primal LP feasibility tolerance).

```
>> feasTol = getCobraSolverParams('LP', 'feasTol');
>> param.boundPrecisionLimit = feasTol*10;
```

`param.tissueSpecificSolver`: the name of the tissue-specific solver to be used to extract the context-specific model (possible options: 'thermoKernel' and 'fastcore'; default: 'thermoKernel').

`param.TolMaxBoundary`: the reaction-flux upper-bound maximum value (default: 1,000).
`param.TolMinBoundary`: the reaction-flux lower-bound minimum value (default: -1,000).
`specificData.rxns2constrain`: table containing the reaction identifier, the updated lower bound, the updated upper bound, a constraint description and any notes such as references or special cases (default: empty).

rxns	Lower bound	Upper bound	Constraint description	Notes
rxn1	0		Close uptake	PMID: ****
rxn2	0.956	1,000	Turnover constraints	PMID: ****

Set coupled reactions (optional)

▲ **CRITICAL** If `param.addCoupledRxns` and the table `specificData.coupledRxns` are empty, this step is not performed.

▲ **CRITICAL** Often, published biochemical data specify metabolites that are known to be produced or metabolized by the biological system. However, a model frequently includes multiple production or degradation pathways. To ensure that the net production or consumption of the metabolite is represented correctly in the model, the concept of coupled reactions can be used³¹. In nongrowing, nondividing cells, such as neurons, these constraints can be used to replace the biomass reaction to set constraints that specify specific biochemical requirements for cell maintenance⁹.

15. (Optional) Unless using default settings, set the coupled reactions on `specificData.coupledRxns`. Coupling constraints are denoted by the matrix inequality $Cv \leq d$, where the matrix $C \in \{-1, 1\}^{k \times n}$ specifies whether a metabolite is produced or consumed and $d \in \mathbb{R}^k$ specifies the net value of the production or consumption of that metabolite. That is, $C_{ij} = -1$ if metabolite i is consumed by reaction j and $C_{ij} = 1$ if metabolite i is produced by reaction j . This constraint ensures that several of the coupled reactions are active in the model, and the total production or consumption is at least the amount given by d_i . However, the exact flux partition of total flux within the set of reactions is not specified. The coupling constraints are added to the draft model by including fields such as the constraint matrix containing reaction coefficients `model.c`, the constraint identifiers `model.ctrs`, the constraint senses `model.dsense` and the constraint right-hand side values `model.d`. Review the log; if `param.printLevel` is greater than zero, the information about the added coupled reactions will be printed. For the i^{th} coupling constraint, when `model.dsense(i) = 'G'` this implies:

$$C_{i,1} \times v_1 + C_{i,2} \times v_2 + \dots + C_{i,n} \times v_n \geq d_i,$$

and when `model.dsense(i) = 'L'` it implies:

$$C_{i,1} \times v_1 + C_{i,2} \times v_2 + \dots + C_{i,n} \times v_n \leq d_i.$$

Within the pipeline, the function `addCOBRAConstraints` adds the specified set of coupling constraints to the model.

◆ TROUBLESHOOTING

Usable variables for Step 15

`param.addCoupledRxns`: logical, should the coupled constraints be added (default: false).

`specificData.coupledRxns`: table containing information about the coupled reactions.

This includes the identifier for the coupling constraint (`couplingConstraintID`), the list of coupled reactions (`coupledRxnsList`), the coefficients of those reactions (`c`, given in the same order), the right-hand side of the constraint (`d`), the constraint sense or the directionality of the constraint (`dsense`) and the reference (default: empty).

couplingConstraintID	coupledRxnsList	c	d	dsense	Reference
cRxn1	crxn1, crxn2	1 1	2.0250	G	PMID: ****
cRxn2	crxn1, crxn3	1 -1	0.2265	G	PMID: ****

Remove inactive reactions

▲ **CRITICAL** Before the extraction process, it is crucial to selectively exclude certain reactions from the draft model, emphasizing the importance of refining the model's accuracy and relevance. This process ensures that subsequent computational analyses are conducted on a streamlined and biologically meaningful metabolic network.

16. (Optional) Identify and list metabolic reactions that should be inactive by a manual curation of literature (`specificData.inactiveReactions`). Designate reactions for removal from the draft model, considering them nonessential (`specificData.rxns2remove.rxns`).

This encompasses reactions such as sink or demand reactions that might not be deemed vital for the model. In the event of a discrepancy between input data (for example, a reaction assigned to be inactive on the basis of manual literature curation but active according to Step 5), specify the preference between curated data or omics data from experiments using the parameter `param.curationOverOmics`.

- Set `param.curationOverOmics = true` to prioritize the manual literature curation
- Set `param.curationOverOmics = false` to prioritize the experimental data

After reaction removal, the feasibility of the model is tested, and, if it fails, several reaction bounds will be relaxed via [relaxedFBA](#).

Review the log; if `param.printLevel` is greater than zero, the number of reactions removed and the reactions that were assigned as inactive and were removed will be printed.

Usable variables for Step 16

`param.curationOverOmics`: logical, indicates whether curated data should take priority over omics data (default: false).

`specificData.inactiveReactions`: list of reactions known to be inactive based on manual literature curation (default: empty).

`specificData.rxns2remove.rxns`: list of reactions to remove from the generic model (default: empty).

Remove inactive genes

▲ **CRITICAL** Building on the previous step (Step 16), selectively removing certain genes from the draft model enhances both the quality and specificity of the metabolic network.

17. Identify and list metabolic genes known to be inactive based on the bibliomics data (`specificData.inactiveGenes`). The genes assigned as inactive in bibliomic or transcriptome data are removed from the draft model by `XomicsToModel`, together with the reactions affected by their deletion (identified using the function [allRxnPerActiveGene](#)). However, any reaction assigned to be active in Step 5 will not be deleted at this step. Furthermore, if removing reactions that correspond to inactive genes compromises the feasibility of a draft model, the bounds of the fewest number of those reactions are relaxed via [relaxedFBA](#), and the rest are removed. To ensure that the model is feasible, the relaxed reactions are added to the list of active reactions. In the event of a discrepancy between the datasets, for example, a gene assigned to be inactive on the basis of the manual literature curation but active on the basis of the proteomics data or transcriptomics data, or both (Step 7), preference will be given on the basis of the value of `param.curationOverOmics` variable.

Review the log; if `param.printLevel` is greater than zero, the deleted reactions and genes will be printed, as well as those that were not deleted.

Usable variables for Step 17

`param.curationOverOmics`: logical, indicates whether curated data should take priority over omics data (default: false).

`specificData.inactiveGenes`: list of Entrez identifiers of genes known to be inactive based on the bibliomics data (default: empty).

Find flux consistent subset

▲ **CRITICAL** `XomicsToModel` finds and extracts the largest subset of the draft model that is flux consistent, that is, each reaction that can carry a nonzero steady-state flux.

18. Unless using default settings, specify the parameters `param.fluxCCmethod` and `param.fluxEpsilon` to indicate the algorithm to be used to identify the flux-consistent subset of the draft model and to determine the minimum flux magnitude considered nonzero. All reactions that are flux inconsistent and all metabolites that exclusively participate in flux-inconsistent reactions are removed by `XomicsToModel`. If a flux-

inconsistent metabolite or reaction is identified as active in Step 5, `XomicsToModel` removes it from the list.

Review the log; if `param.printLevel` is greater than zero, `XomicsToModel` prints the dimensions of the stoichiometric matrix before and after checking flux consistency, as well as the constraints for flux consistent reactions and metabolites.

In addition, different Boolean vectors indicating the flux consistent and inconsistent metabolites and reactions are added to the draft model by `XomicsToModel`.

Usable variables for Step 18

`param.fluxCCmethod`: string with the name of the algorithm to be used for the flux consistency check (possible options: 'swiftcc', 'fastcc' or 'dc', default: 'fastcc').

`param.fluxEpsilon`: minimum nonzero flux value accepted for tolerance (default: primal feasibility tolerance).

```
>> feasTol = getCobraSolverParams('LP', 'feasTol');
>> param.fluxEpsilon = feasTol * 10;
```

Find thermodynamically consistent subset (optional)

▲ **CRITICAL** In this step, `XomicsToModel` finds the largest thermodynamically consistent subset. The performance of `thermoKernel` is accelerated if provided with a thermodynamically consistent model before extracting a subset. `XomicsToModel` finds the largest thermodynamically consistent subset using `findThermoConsistentFluxSubset`, which requires `model.S`, `model.lb` and `model.ub` as essential inputs.

19. (Optional) Unless using default settings, specify `param.tissueSpecificSolver = 'thermoKernel'` as the extraction algorithm. Specify the metabolites and reactions that are flux consistent using `model.fluxConsistentMetBool` and `model.fluxConsistentRxnBool` to accelerate this step.

All reactions that are not thermodynamically consistent are removed from the draft model by `XomicsToModel`. All metabolites that are exclusively involved in thermodynamically inconsistent reactions are also removed from the draft model by `XomicsToModel`. If a thermodynamically inconsistent metabolite or reaction is identified as active in `specificData.presentMetabolites` or in the list of active reactions from Step 5, it is removed from the list by `XomicsToModel`. In addition, all orphan genes are removed from the model and the reaction–gene–matrix is regenerated. Boolean vectors indicating thermodynamically consistent metabolites and reactions are added to the draft model, `model.thermoFluxConsistentMetBool` and `model.thermoFluxConsistentRxnBool`, respectively.

Review the log; if `param.printLevel` is greater than zero, `XomicsToModel` prints the dimensions of the stoichiometric matrix before and after checking thermodynamic consistency, as well as the constraints for thermodynamic consistent reactions and metabolites.

To identify a thermodynamically feasible flux, examine the net flux vector where each nonzero entry indicates a sign opposite the change in chemical potential for the corresponding reaction. A model is deemed thermodynamically consistent if each of its reactions enables a nonzero thermodynamically feasible flux, thereby meeting the criteria for consistency. This is outlined by the following mathematical definition:

Usable variables for Step 19

`param.tissueSpecificSolver`: the name of the tissue-specific solver to be used to extract the context-specific model (possible options: 'thermoKernel' and 'fastcore'; default: 'thermoKernel').

`param.thermoFluxEpsilon`: flux epsilon used in `thermoKernel` (default: solver primal feasibility tolerance×10).

`param.iterationMethod`: enables different iteration methods to be employed when exploring the thermodynamically consistent set.

Identify active reactions from genes

▲ **CRITICAL** The list of active reactions up to this point has been based on the subset of user-specified context-specific reactions that are determined to be flux consistent by `XomicsToModel`. The relationship between context-specific genes and metabolic reactions is established by `XomicsToModel` in this step. Each identified reaction is added by `XomicsToModel` to the final list of active reactions used to extract the context-specific model.

20. Decide how `XomicsToModel` determines the gene–reaction relationship and specify your preference using `param.activeGenesApproach`.

- `'oneRxnPerActiveGene'`: at least one reaction per active gene is included (default)
- `'allRxnPerActiveGene'`: find a list of reactions whose rates are affected by the deletion of active genes and include them all as active reactions

Changing this parameter has a large effect on the size of the extracted models. Per active gene, `oneRxnPerActiveGene` adds at least one corresponding reaction to the extracted model, whereas `allRxnPerActiveGene` adds all corresponding reactions, where the presence of the gene product is essential for nonzero flux through each of the corresponding reactions. Experience with the generation of whole-body metabolic models³⁵ and extraction of a neuronal model⁹ from Recon3D supports the use of `oneRxnPerActiveGene` as the default.

Usable variable for Step 20

`param.activeGenesApproach`: string with the name of the active genes approach will be used (possible options: `'oneRxnPerActiveGene'`, `'allRxnPerActiveGene'`; default: `'oneRxnPerActiveGene'`).

Model extraction

▲ **CRITICAL** In this step, given a draft model, a specific model is extracted using the function `createTissueSpecificModel`. The `createTissueSpecificModel` function is already enabled to extract context-specific models using a suite of model extraction algorithms (GIMME²¹, iMAT²², MBA⁷, INIT⁵⁰, mCADRE²³ and FASTCORE¹²) and has been described previously (step 40 of a previous protocol³⁴). Each of these algorithms extracts a flux-consistent context-specific model. In `XomicsToModel`, the default model extraction algorithm is `thermoKernel`, because it enables extraction of a context-specific model that is stoichiometrically consistent, flux-consistent and thermodynamically consistent, now accessible via the `createTissueSpecificModel`. The implementation of thermodynamic constraints during the model extraction process enables the integration of context-specific data on each metabolite, which is not directly integrable with the other aforementioned model extraction algorithms.

21. Set the `param.tissueSpecificSolver` parameter to guide on the choice of solver: the default option `'thermoKernel'` extracts a thermodynamically consistent context-specific model, ensuring a robust foundation for further analyses. Alternatively, `'fastCore'` is for extracting a minimal set of reactions supporting optimal flux solutions, emphasizing the preservation of essential functionalities.

Usable variable for Step 21

`param.tissueSpecificSolver`: the name of the tissue-specific solver to be used to extract the context-specific model (possible options: `'thermoKernel'` and `'fastCore'`; default: `'thermoKernel'`).

thermoKernel (optional)

▲ **CRITICAL** `thermoKernel` is an algorithm designed to extract a thermodynamically consistent context-specific model. It is important that the weightings on metabolites and weightings on reactions (derived from experimental data) are appropriately adjusted by the user to ensure accurate and context-specific model outcomes.

22. (Optional) To assign weights to metabolites, on the basis of quantitative data or through manual curation, set the weights of the metabolites in `specificData.presentMetabolites.weights` accordingly ('Experimental design' section).

Likewise, for assigning weights to absent metabolites, set `specificData.absentMetabolites.weights` to values < 0 . The negative weights promote inclusion of metabolites and reactions in a specific model, positive weights do the opposite and a zero weight does not penalize for exclusion or inclusion.

By default, `XomicsToModel` sets a fixed penalty for the inclusion of any metabolite or reaction not in a core set. An exception is that `XomicsToModel` sets a weight of zero for any metabolite ranked in the top 100 for metabolite connectivity, mostly nonspecific cofactors. This is to avoid the exclusion or inclusion metabolites, such as ATP, that participate in any context-specific network. Therefore, the objective of integrating metabolite weights is to promote the exclusion or inclusion of metabolites that participate in a small number of relatively pathway-specific reactions.

Add a penalty for reactions: provide gene expression data (Step 7) to set the fixed penalty equal to the median of the log-transformed reaction expression values in `model.expressionRxns`. In the absence of expression data, the fixed penalty is set to one. Use `param.activeGenesApproach = 'oneRxnPerActiveGene'` (Step 20) to specify that all metabolites and reactions in the core sets are assigned a weight equal to the negative of the median of the log of the reaction expression value `model.expressionRxns` or -1 if `model.expressionRxns` is not provided.

If `model.expressionRxns` is provided by the user, but a reaction is missing gene expression data, `XomicsToModel` assumes that the corresponding reaction has a weight of equal to the negative of the median of the log of the reaction expression values `model.expressionRxns`.

Context-specificity and thermodynamic consistency: a model is context-specific and thermodynamically flux consistent if it meets two criteria:

- Thermodynamic flux consistency: each reaction admits a nonzero thermodynamically feasible flux
- Context-specific enrichment: the model includes metabolites and reactions relevant to a particular context while excluding those that are not relevant to that context

Evaluate the extracted model: an extracted model should be considered a draft model until it has been compared with the expected results of model extraction. For example, if `thermoKernel` was selected as the model extraction algorithm, then all reactions should admit a thermodynamically feasible flux. Furthermore, the draft model should be checked to see if it contains the anticipated metabolites and reactions. If not, then the adjustment of parameters or input data may be warranted.

Review the log; if `param.printLevel` is greater than zero, the size of the stoichiometric matrix, the metabolites and reactions removed by the extraction algorithm, and the core metabolites and reactions removed by the extraction solver are all printed.

- If greater than 1, the progress of each outer iterations (equation (6)) of the `thermoKernel` algorithm is printed, including the fraction of thermodynamically consistent reactions and metabolites that have been added to the extracted model at each iteration and their overlap with the incentivized and penalized reactions and metabolites
- If greater than 2, the progress of each inner iteration (equation (5)) of the `thermoKernel` algorithm is printed, including the number of nonzero thermodynamically feasible fluxes identified in each iteration, before and after checking for thermodynamic feasibility, thus enabling the user to monitor the progress of the model extraction

◆ TROUBLESHOOTING

Usable variables for Step 22

`param.activeGenesApproach`: string with the name of the active genes approach will be used (possible options: `'oneRxnPerActiveGene'`, `'allRxnPerActiveGene'`; default: `'oneRxnPerActiveGene'`).

`param.fluxEpsilon`: minimum nonzero flux value accepted for tolerance (default: primal feasibility tolerance).

Protocol

```
>> feasTol = getCobraSolverParams('LP', 'feasTol');  
>> param.fluxEpsilon = feasTol * 10;
```

`param.thermoFluxEpsilon`: minimal flux considered thermodynamically consistent by `thermoKernel` (default: primal feasibility tolerance $\times 10$).

`param.tissueSpecificSolver`: the name of the tissue-specific solver to be used to extract the context-specific model (possible options: 'thermoKernel' and 'fastCore'; default: 'thermoKernel').

`param.weightsFromOmics`: true for gene weights to be assigned on the basis of the omics data (default: true).

`specificData.absentMetabolites`: list of metabolites confirmed to be absent on the basis of the bibliomics data and their corresponding weights.

`specificData.presentMetabolites`: list of metabolites confirmed to be present on the basis of the bibliomics data and their corresponding weights.

Final adjustments

▲ **CRITICAL** In the last step, `XomicsToModel` checks the flux and thermodynamic consistency of the model and adds any specific data and the parameters used in the extraction of the specific model to the model structure, that is, `model.XomicsToModelSpecificData` and `model.XomicsToModelParam`, respectively. If a field with metabolic reaction formulas did not already exist in the draft model, `XomicsToModel` adds it using the [printRxnFormula](#) function. Also, `XomicsToModel` adds two vectors that indicate which reactions were specified to be active or inactive (`model.activeInactiveRxn`), as well as which metabolites were specified to be present or absent (`model.presentAbsentMet`). Finally, `XomicsToModel` reorders the fields of the context-specific model into a standard format.

23. Review the log entries to confirm the model's dimensions at various steps in the `XomicsToModel` process. For example, the effect of major steps in the pipeline are evident by the relative increase or decrease in size of the draft model; therefore, unexpected changes in model size should be investigated in the corresponding step to ensure that the input data are correctly specified and the technical parameters are appropriate to the context.

If `param.printLevel` was set to > 0 , the algorithm will display all outputs from the `XomicsToModel` function. This detailed output can be useful, especially in identifying steps where reactions were removed owing to a lack of flux consistency, stoichiometric consistency or thermodynamic consistency. Such insights serve as cues for potential manual curation, allowing users to address specific reactions that might have been excluded but could be of interest. Moreover, having a detailed log facilitates traceability, enabling users to understand the evolution of the model throughout the process. It serves as a documentation tool, aiding in debugging, and provides a comprehensive record of the decisions made during model construction, offering transparency and reproducibility in the analysis.

Troubleshooting

Troubleshooting advice can be found in [Table 2](#).

Table 2 | Troubleshooting

Step	Problem	Possible reason	Solution
All	Unexpected behavior within a key step	Erroneous or inconsistent input data. Exceptional use case for <code>XomicsToModel</code> not observed previously	Set <code>param.debug = true</code> ; and run <code>XomicsToModel</code> again. All results generated before a particular step are saved in a MATLAB workspace file (.mat) corresponding to that step. Therefore, in the case of unexpected results, rather than running the entire pipeline again, one can load the MATLAB workspace file at the desired step then restart <code>XomicsToModel</code> in the MATLAB debugging mode as required

Table 2 (continued) | Troubleshooting

Step	Problem	Possible reason	Solution
2	Infeasible generic input model	Flux or stoichiometric inconsistency, or because its bounds are overconstrained	Either the model or the included information needs to be modified. The fastGapFill function can add missing reactions to make the draft model feasible. Reactions can be added using the addReaction function after additional manual literature curation. If it is overconstrained, relaxedFBA relaxes the minimum number of reaction bounds in the draft model to make it feasible
4	RXN1 in specificData.rxns2add. rxns requires a reaction formula in specificData.rxns2add.rxnFormulas	New metabolic reactions in the draft model require the addition of the metabolic reaction formula	Place the metabolic reaction formula in the specificData. rxns2add.rxnFormulas for each reaction in the following format: 'met1 + 2 met2 → met3 + met4' where the type of reaction is indicated with the proper characters (forward (→); reverse (←); bidirectional (↔))
14	Lower bounds greater than upper bounds	A lower bound (lb) of a reaction cannot be greater than an upper bound (ub) $lb \leq ub$	Check the bounds included in the options variable including: specificData.rxns2add, param.TolMinBoundary, param. TolMaxBoundary, specificData.rxns2constrain and specificData. exoMet described in Steps 4, 6, 12 and 13
15	Missing fields in the table	To add the coupled reactions, no element of the table specificData. coupledRxns can be left empty	Fill in the empty cells with the missing information as described in Step 15
22	Extracted model contains reactions that do not admit a thermodynamically feasible net flux	The objective term $\beta^T(p + q)$, which minimizes the absolute value of internal reaction flux $z \in \mathbb{R}^n$, which promotes thermodynamic consistency, is too small	Increase the value of the tradeoff parameter (model.beta) on minimization of one-norm of internal fluxes. The value of β should be such that $\beta^T(p + q) + g^T z _0 + h^T s _0 + f^T r _0$ at the optimum of each iteration within the thermoKernel algorithm
22	Extracted model does not contain the anticipated metabolites	Anticipated metabolites must be given negative weights	Check that anticipated metabolite identifiers in specificData. presentMetabolites.mets match those of the metabolites in the model. If weights are provided, ensure the sign of the weights are all negative in specificData.presentMetabolites. weights
		Anticipated metabolites may not be part of the flux (or thermodynamically consistent subset)	Check whether the anticipated metabolites were eliminated before the model extraction step (see the Supplementary Information for the description of function debugXomicsToModel)
		Negative weighs on metabolites may be too small a magnitude relative to weights on reactions or strictly positive parameter β , which optimizes toward a thermodynamically feasible flux	Check if increasing the magnitude of the negative weights on metabolites (specificData.presentMetabolites.weights) increases the fraction of anticipated metabolites included in the extracted model
22	Extracted model does not contain the anticipated reactions	Anticipated reactions must be either core reactions or incentivized to be present by corresponding transcriptomic or proteomic data	Check that: (1) the anticipated reactions are core reactions specified in specificData.activeReactions by an identifier that matches one in the generic model; (2) the gene expression specificData.transcriptomicData or proteomic data specificData. proteomics are expressed above the corresponding thresholds; (3) the gene–protein–reaction association is correct for that reaction in the generic model, model.grRules; and (4) for the reactions to be constrained, specificData.rxns2constrain contains bounds of sufficient magnitude (that is, at least $10 \times$ the cutoff for nonzero flux, which by default is 1×10^{-5} and therefore typically of magnitude 1×10^{-4} or greater)
		Anticipated reactions may not be part of the flux (or thermodynamically consistent subset). Negative weighs on reactions may be too small a magnitude relative to weights on metabolites or strictly positive parameter β , which optimizes toward a thermodynamically feasible flux	Check that: (1) the anticipated reactions were eliminated before the model extraction step (see the Supplementary Information for the description of function debugXomicsToModel); (2) if gene expression or proteomic data are provided for anticipated reactions, the gene expression specificData. transcriptomicData or proteomic data specificData.proteomics for the corresponding genes results in negative reaction weights in rxnWeights that are passed to thermoKernelWeights; and (3) increasing the magnitude of the expression or proteomic data increases the fraction of anticipated reactions included in the extracted model

Table 2 (continued) | Troubleshooting

Step	Problem	Possible reason	Solution
22	Extracted model contains metabolites or reactions that were not anticipated	Metabolites or reactions may need to be added by the model extraction algorithm to ensure that other core metabolites and reactions are (thermodynamically) flux consistent	If these metabolites or reactions should not be present in the generated model, then review literature and add additional pathways that enable the model to include anticipated metabolites and reactions in the newly added functional pathways
		Positive weight on metabolites or reactions is not of a large enough magnitude relative to β	Check if increasing the magnitude of the positive weights on absent metabolites (<code>specificData.absentMetabolites.weights</code>) reduces the number of these unwanted metabolites included in the extracted model, or check if decreasing the magnitude of the weights on reactions (derived from <code>specificData.transcriptomicData.expVal</code> and/or <code>specificData.proteomicData.expVal</code>) reduces the number of unwanted reactions included in the extracted model
22	Extracted model does not admit sufficient flux through a specified biomass reaction	Essential reactions or precursor metabolites are missing, or the imposed flux bounds on the biomass reaction restrict feasible flux	Ensure <code>param.biomassRxn</code> is set to the correct reaction abbreviation. Add a sufficient lower bound on the biomass reaction flux as an additional constraint that is compatible with a steady-state flux in the input generic model <code>specificData.rxns2constrain</code>

Timing

Step 1, data preparation: ~5 s
 Step 2, generic model check: ~1 s
 Step 3, set objective function: 1 s
 Step 4, add missing reactions: ~30 s
 Step 5, identify core metabolites and reactions: ~10 s
 Step 6, set limit bounds: ~1 s
 Step 7, identify active genes: ~100 s
 Step 8, close ions: ~1 s
 Step 9, close exchange reactions: ~1 s
 Step 10, close sink and demand reactions: ~10 s
 Step 11, set metabolic constraints: ~10 s
 Step 12, cell culture data: ~10 s
 Step 13, metabolomics: ~10 s
 Step 14, add custom constraints: ~10 s
 Step 15, set coupled reactions: 10 s
 Step 16, remove inactive reactions: 10 s
 Step 17, remove inactive genes: ~100 s
 Step 18, find flux-consistent subset: ~100 s
 Step 19, find thermodynamically consistent subset: ~10 s
 Step 20, identify active reactions from genes: ~10 s
 Step 21, model extraction: ~100 s
 Step 22, thermoKernel: ~100 s
 Step 23, final adjustments: ~1 s

Anticipated results

This section follows the steps of the procedure. The parameters used and modified are printed in the diary described in Step 2. Any results or anomalies are also printed in this diary.

1. The parameters that will be used by the `XomicsToModel` pipeline are printed, including the default values for options that have not been declared.

Table 3 | Recon3D²⁴ stoichiometric consistency summary

Mets (metabolites)	Rxns	Analysis
5,835	10,600	Totals
0	1,809	Heuristically external
5,835	8,791	Heuristically internal
5,835	8,791	... of which are stoichiometrically consistent
0	0	... of which are stoichiometrically inconsistent
0	0	... of which are of unknown consistency
0	0	Heuristically internal and stoichiometrically inconsistent or unknown consistency
0	0	... of which are elementally imbalanced (inclusively involved metabolite)
0	0	... of which are elementally imbalanced (exclusively involved metabolite)
5,835	8,791	Confirmed stoichiometrically consistent by leak and/or siphon testing

- Any modifications made at this step are indicated; for example, if the generic model contains `DM_atp_c_reaction` it is renamed to `ATPM` as it represents ATP maintenance and to ensure that it is not closed together with standard demand reactions for the model extraction with `thermoKernel`. Furthermore, if the compartment representing mitochondrial intramembrane space ([i]) is removed (if `thermoKernel` is chosen), the modified reactions are also printed out.
- The output of this set is a message including the specified linear objective function or a lack thereof.
- The reactions and formulas added in this step are printed, together with the results of the stoichiometric consistency test. An example of the stoichiometric consistency test for Recon3D²⁴ can be seen in Table 3.
- The output of this set is a set of present metabolites and reactions identified on the basis of bibliomic, metabolomic and cell culture data.
- The output of this set is a draft model with the updated maximum upper bounds and minimum lower bounds.
- The output of this set is a set of active genes identified using bibliomic, transcriptomic and proteomic data. In addition, a message is displayed indicating the number of genes that have no expression information owing to a lack of transcriptomic data.
- The output of this set is a draft model in which the upper and lower bounds of ion exchange reactions, such as calcium, potassium or sodium, among others, are set to zero.
- The output of this set is a draft model where all the lower bounds of exchange reactions are set to zero, except for those exchange reactions corresponding to metabolites specified to be present in the environment, for example, fresh media metabolites, which are assumed to be possible to be taken up from the media. Exchange reaction upper bounds are unchanged. A message is shown reporting model statistics such as the stoichiometric matrix dimensions, the number of closed reactions, the number of all exchange reactions and the number of exchange reactions in the active reactions.
- The sink and demand reactions of the draft model are set to zero and a message is shown reporting the number of closed noncore sink and/or demand reactions and the method used in `param.nonCoreSinksDemands`.
- The output of this set is a draft model with added metabolomic constraints from the environment, for example, cell culture exometabolomic constraints or both. Compared with the draft model from the previous step, which usually derives generic upper and lower bounds from the generic input model, the addition of metabolomic constraints from the environment results in a subset of exchange reactions with quantitatively specified upper and lower bounds. In addition, a message is displayed reporting the number of exometabolomic and media uptake constraints. It includes exometabolomic data statistics, the fitting procedure for each exometabolomic reaction, the relaxed reactions and the summary of the fitting.

12. The output of this set is a draft model with custom constraints based on bibliomic data. If `thermoKernel` is being used, demand reactions included in `specificData.rxns2constrain`, that is, with prefix `DM_`, will be ignored, and the message indicating this is printed. Furthermore, a message is given showing a list of reaction identifiers that cannot be constrained because they are not present in the model; the number of constrained reactions; and whether or not the model was still feasible after the addition of custom constraints.
13. The output of this set is a draft model that includes new fields that correspond to the coupled reactions. A message is shown including the details of the coupled constraints imposed on the draft model.
14. The output of this set is a draft model, with inactive reactions removed (based on bibliomic data). A message including the number of reactions that were removed from the model, a list of reactions that were not present in the draft model before the reaction removal, as well as the results of the model feasibility check is printed.
15. The output of this set is a draft model with inactive genes removed and a detailed description of the number of genes that were removed or missing in draft model before this step. Based on the value of `param.curationOverOmics`, a message is given about the genes that were kept or removed from the manually curated set of inactive genes due to the conflict with the provided omics data (measured as active in transcriptomics and/or proteomics). Furthermore, information is given about the number of genes that were not removed because their removal would lead to an infeasible model. Lastly, model feasibility check results are given, including any reaction that was relaxed if the model was not feasible.
16. The output of this set is a draft model with flux-consistent reactions and metabolites printed, along with a report containing model statistics such as stoichiometric matrix dimensions, flux-consistent and flux-inconsistent reactions and metabolites.
17. The output of this set is a thermodynamically consistent draft model. In addition, a message is given with the parameters used to identify the thermodynamically consistent subset. It is followed by a detailed report from the cardinality optimization. Lastly, the stoichiometric matrix dimensions, thermodynamically consistent and inconsistent reactions and metabolites are printed out.
18. The output of this set is a set of reactions that must be active on the basis of the active genes.
19. The output of this set is an extracted genome-scale metabolic model.
 - If `fastCore` is selected as the extraction algorithm, a message is printed out with the extraction summary and the size of the new model
 - If `thermoKernel` is selected as the extraction algorithm, a message is printed out with the parameters used in the `thermoKernel` function, as well as `optCardThermo` and `optimizeCardinality` functions called by the `thermoKernel` function, an extraction summary and statistics for the new model

Two confusion matrices are generated that compare specified and actual metabolites and reactions in an extracted model (Fig. 4). This comparison validates the accuracy and completeness of the extraction process, ensuring that the model faithfully represents the biological system. An accurate validation is crucial for reliable analyses and simulations based on the model. A figure with comparison of weights on active and inactive reactions and metabolites is also generated (Fig. 5). This is useful for assessing if the relative weights on present and absent metabolites, as well as active and inactive reactions, passed as parameters to `thermoKernel`, are sufficient to incentivize and disincentivize the presence or absence of metabolites and reactions in the generated model. As the generated model is a compromise between the weights on metabolites and reactions, one may find that several metabolites and reactions may be omitted at the expense of inclusion of other reactions and metabolites, respectively, depending on the input weights to `thermoKernel`. The user can modify the weights in `param.weightsFromOmics`, `presentMetabolites.weights` or `absentMetabolites.weights` to control the inclusion or exclusion of specific metabolites and reactions.

20. Final adjustments: the output of this step is the final context-specific flux-consistent genome-scale metabolic model. If `thermoKernel` was chosen, then the model is also




Fig. 4 | Comparison of specified and actual metabolites and reactions in an extracted model. The confusion matrices between the reactions and metabolites specified to be active and inactive versus output model reactions and metabolites that are actually active and inactive, obtained while generating a context-specific model of a dopaminergic neuronal metabolism⁹. The end result may be a tradeoff between incentives to be active (negative weights), incentives to be inactive (positive weights) and ambivalence (zero weights). When a reaction is incentivized to be present, but the corresponding metabolites are incentivized to be inactive, the optimal solution at each iteration of thermoKernel will depend on the relative weights on the reactions versus metabolites. This plot is automatically generated by XomicsToModel if `param.plotThermoKernelStats = 1`, when `param.tissueSpecificSolver = 'thermoKernel'`.

thermodynamically flux consistent. With all input data fully prepared, the algorithmic completion of the pipeline takes ~10 min; however, manual review of intermediate results may also be required, for example, when inconsistent input data lead to an infeasible model. Previous applications of the XomicsToModel pipeline illustrate the type of results that can be anticipated. For instance, dopaminergic neuronal models generated with the Protocol reproduced experimentally observed metabolite exchange rates and highlighted bioenergetic vulnerabilities relevant to Parkinson's disease^{9,25}. Multiomics integration studies demonstrated that XomicsToModel yields reproducible, condition-specific models that capture context-dependent metabolic rewiring^{26,27}. Furthermore, its use in isotope-labeled flux studies shows the potential of the Protocol to provide mechanistic insight into immune metabolism and quantitative flux predictions^{28,29}. Together, these results suggest that users of XomicsToModel can expect to obtain thermodynamically consistent, biologically faithful models that are immediately applicable to hypothesis generation and disease-oriented research.




Fig. 5 | Comparison of weights on active and inactive reactions and metabolites. Almost all reactions incentivized to be active (negative weights) are active, with a small proportion inactive that should be active (reactions omitted, light blue), obtained while generating a context-specific model of a dopaminergic neuronal metabolism⁹. Reactions, or metabolites, with zero weight (0) are ambivalent with respect to inclusion in the model or not. A small positive (+ve) penalty penalizes inclusion of remaining reactions and metabolites to optimize for a minimal model size. All metabolites incentivized to be active (negative (-ve) weights) are active, with none inactive that should be active (metabolites omitted, green in legend, no bar in plot). This plot is automatically generated by XomicsToModel if `param.plotThermoKernelStats = 1`, when `param.tissueSpecificSolver = 'thermoKernel'`.

Data availability

Constraint-based reconstruction, modeling and analysis was implemented in MATLAB (MathWorks). Open-source computer code enabling the reproduction of all computational steps is available within the COBRA Toolbox³⁴ version 3.4+. This includes code for the model generation `XomicsToModel.m`, thermodynamically feasible model extraction `thermoKernel.m`, debugging the performance of `XomicsToModel` `debugXomicsToModel.m`, the comparison of multiple models `compareXomicsToModel.m` and for the maximization of flux entropy, with the option of quadratic penalization of deviation from measured experimental fluxes `entropicFBA.m`; it is available via GitHub at <https://github.com/opencobra/cobratoolbox/tree/master/src/dataIntegration/XomicsToModel> (ref. 51). An executable narrative tutorial demonstrating the use of the XomicsToModel pipeline is also available, `tutorial_XomicsToModel.mlx`, with the example data for generating a dopaminergic neuronal metabolic model; it is available via GitHub at https://github.com/opencobra/COBRA.tutorials/blob/master/dataIntegration/XomicsToModel/tutorial_XomicsToModel.mlx (ref. 52). An HTML version of this tutorial is accessible via GitHub at https://github.com/opencobra/COBRA.tutorials/blob/master/dataIntegration/XomicsToModel/tutorial_XomicsToModel.html. Linear and quadratic optimization problems may be solved with open source solvers, for example, **GLPK**, or using an industrial solver, such as Gurobi 9.1 (Gurobi). Nonlinear convex optimization problems, for example, entropic optimization, may be solved with the open source solver primal-dual interior method for convex objectives **PDCO** implemented in MATLAB (MathWorks) or with an industrial solver, using the exponential cone solver within Mosek 10.0.30 (Mosek ApS). Note that PDCO may require specialist numerical optimization parameters to be tailored to each model, whereas Mosek is more robust with respect to the numerical characteristics of input models. Each of the aforementioned solvers are interfaced with the COBRA Toolbox and are suited for models that are not numerically ill-scaled.

Received: 2 June 2023; Accepted: 26 September 2025;

Published online: 1 December 2025

References

- Palsson, B. Ø. *Systems Biology: Constraint-Based Reconstruction and Analysis*. (Cambridge Univ. Press, 2015).
- Thiele, I. & Palsson, B. Ø. A protocol for generating a high-quality genome-scale metabolic reconstruction. *Nat. Protoc.* **5**, 93–121 (2010).
- Machado, D. et al. Fast automated reconstruction of genome-scale metabolic models for microbial species and communities. *Nucleic Acids Res.* **46**, 7542–7553 (2018).
- Norsigian, C. J. et al. A workflow for generating multi-strain genome-scale metabolic models of prokaryotes. *Nat. Protoc.* **15**, 1–14 (2020).
- Wang, H. et al. Genome-scale metabolic network reconstruction of model animals as a platform for translational research. *Proc. Natl Acad. Sci. USA* <https://doi.org/10.1073/pnas.2102344118> (2021).
- Opdam, S. et al. A systematic evaluation of methods for tailoring genome-scale metabolic models. *Cell Syst.* **4**, 318–329.e6 (2017).
- Jerby, L., Shlomi, T. & Ruppin, E. Computational reconstruction of tissue-specific metabolic models: application to human liver metabolism. *Mol. Syst. Biol.* <https://doi.org/10.1038/msb.2010.56> (2010).
- Vinnakota, K. C. et al. Network modeling of liver metabolism to predict plasma metabolite changes during short-term fasting in the laboratory rat. *Front. Physiol.* **10**, 161 (2019).
- Preciat, G. et al. Mechanistic model-driven exometabolomic characterisation of human dopaminergic neuronal metabolism. Preprint at *bioRxiv* <https://doi.org/10.1101/2021.06.30.450562> (2021).
- Haraldsdóttir, H. S. et al. CHRR: coordinate hit-and-run with rounding for uniform sampling of constraint-based models. *Bioinformatics* **33**, 1741–1743 (2017).
- Fleming, R. M. T. & Thiele, I. Mass conserved elementary kinetics is sufficient for the existence of a non-equilibrium steady state concentration. *J. Theor. Biol.* **314**, 173–181 (2012).
- Vlassis, N., Pacheco, M. P. & Sauter, T. Fast reconstruction of compact context-specific metabolic network models. *PLoS Comput. Biol.* **10**, e1003424 (2014).
- Fleming, R. M. T. et al. A variational principle for computing nonequilibrium fluxes and potentials in genome-scale biochemical networks. *J. Theor. Biol.* **292**, 71–77 (2011).
- Orth, J. D., Thiele, I. & Palsson, B. Ø. What is flux balance analysis? *Nat. Biotechnol.* **28**, 245–248 (2010).
- Thiele, I., Vlassis, N. & Fleming, R. M. T. fastGapFill: efficient gap filling in metabolic networks. *Bioinformatics* **30**, 2529–2531 (2014).
- Beard, D. A., Liang, S.-D. & Qian, H. Energy balance for analysis of complex metabolic networks. *Biophys. J.* **83**, 79–86 (2002).
- Henry, C. S., Broadbelt, L. J. & Hatzimanikatis, V. Thermodynamics-based metabolic flux analysis. *Biophys. J.* **92**, 1792–1805 (2007).
- Dill, K. & Bromberg, S. *Molecular Driving Forces Statistical Thermodynamics in Biology, Chemistry, Physics, and Nanoscience* (Garland Science, 2011).
- Boyd, S. P. & Vandenberghe, L. *Convex Optimization* (Cambridge Univ. Press, 2004).
- Desouki, A. A. et al. CycleFreeFlux: efficient removal of thermodynamically infeasible loops from flux distributions. *Bioinformatics* **31**, 2159–2165 (2015).
- Becker, S. A. & Palsson, B. Ø. Context-specific metabolic networks are consistent with experiments. *PLoS Comput. Biol.* **4**, e1000082 (2008).
- Zur, H., Ruppin, E. & Shlomi, T. iMAT: an integrative metabolic analysis tool. *Bioinformatics* **26**, 3140–3142 (2010).
- Wang, Y., Eddy, J. A. & Price, N. D. Reconstruction of Genome-Scale Metabolic Models for 126 Human Tissues Using mCADRE. *BMC Syst. Biol.* **6**, 153 (2012).
- Brunk, E. et al. Recon3D enables a three-dimensional view of gene variation in human metabolism. *Nat. Biotechnol.* **36**, 272 (2018).
- Luo, X. et al. Constraint-based modeling of bioenergetic differences between synaptic and non-synaptic components of human dopaminergic neurons. *Front. Comput. Neurosci.* **19**, 1594330 (2025).
- Luo, X. et al. Identification of metabolites reproducibly associated with Parkinson's disease via meta-analysis and computational modelling. *NPJ Parkinsons Dis.* **10**, 126 (2024).
- Zagare, A. et al. Omics data integration suggests a potential idiopathic Parkinson's disease signature. *Commun. Biol.* **6**, 1179 (2023).
- Fleming, R. M. T. et al. Conserved moiety fluxomics. Preprint at *bioRxiv* <https://doi.org/10.1101/2024.11.21.624666> (2024).

29. Huang, L. et al. fluxTrAM: integration of tracer-based metabolomics data into atomically resolved genome-scale metabolic networks for metabolic flux analysis. Preprint at *bioRxiv* <https://doi.org/10.1101/2024.11.26.625485> (2024).
30. Kümmel, A., Panke, S. & Heinemann, M. Putative regulatory sites unraveled by network-embedded thermodynamic analysis of metabolome data. *Mol. Syst. Biol.* **2**, 2006.0034 (2006).
31. Thiele, I. et al. Functional characterization of alternate optimal solutions of *Escherichia coli*'s transcriptional and translational machinery. *Biophys. J.* **98**, 2072–2081 (2010).
32. Medlock, G. L., Moutinho, T. J. & Papin, J. A. Medusa: software to build and analyze ensembles of genome-scale metabolic network reconstructions. *PLoS Comput. Biol.* **16**, e1007847 (2020).
33. Balestrino, R. & Schapira, A. H. V. Parkinson disease. *Eur. J. Neurol.* **27**, 27–42 (2020).
34. Heirendt, L. et al. Creation and analysis of biochemical constraint-based models using the COBRA toolbox v3.0. *Nat. Protoc.* **14**, 639 (2019).
35. Thiele, I. et al. Personalized whole-body models integrate metabolism, physiology, and the gut microbiome. *Mol. Syst. Biol.* **16**, e8982 (2020).
36. Pacheco, M. P. et al. scFASTCORMICS: a contextualization algorithm to reconstruct metabolic multi-cell population models from single-cell RNAseq data. *Metabolites* **12**, 1211 (2022).
37. Vieira, V., Ferreira, J. & Rocha, M. A pipeline for the reconstruction and evaluation of context-specific human metabolic models at a large-scale. *PLoS Comput. Biol.* **18**, e1009294 (2022).
38. Estévez, S. R. & Nikoloski, Z. Generalized framework for context-specific metabolic model extraction methods. *Front. Plant Sci.* **5**, 491 (2014).
39. Gopalakrishnan, S. et al. Guidelines for extracting biologically relevant context-specific metabolic models using gene expression data. *Metab. Eng.* **75**, 181–191 (2023).
40. Tefagh, M. & Boyd, S. P. SWIFTCORE: a tool for the context-specific reconstruction of genome-scale metabolic networks. *BMC Bioinform.* **21**, 140 (2020).
41. Akbari, A. et al. The quantitative metabolome is shaped by abiotic constraints. *Nat. Commun.* **12**, 3178 (2021).
42. Mayr, E. *What Makes Biology Unique?: Considerations on the Autonomy of a Scientific Discipline* (Cambridge Univ. Press, 2007).
43. Haraldsdóttir, H. S., Thiele, I. & Fleming, R. M. T. Quantitative assignment of reaction directionality in a multicompartamental human metabolic reconstruction. *Biophys. J.* **102**, 1703–1711 (2012).
44. Sahin, A., Weilandt, D. R. & Hatzimanikatis, A. Optimal enzyme utilization suggests that concentrations and thermodynamics determine binding mechanisms and enzyme saturations. *Nat. Commun.* **14**, 2618 (2023).
45. Nath, S. S. Evaluation of thermodynamic consistency of kinetic parameters in cyclic enzyme-catalyzed reaction networks. *Chem. Phys. Lett.* **804**, 139890 (2022).
46. Kenney, I. M. & Beckstein, O. Thermodynamically consistent determination of free energies and rates in kinetic cycle models. *Biophys. Rep.* **3**, 100120 (2023).
47. Gevorgyan, A., Poolman, M. G. & Fell, D. A. Detection of stoichiometric inconsistencies in biomolecular models. *Bioinformatics* **24**, 2245–2251 (2008).
48. Love, M. I., Huber, W. & Anders, S. Moderated estimation of fold change and dispersion for RNA-seq data with DESeq2. *Genome Biol.* **15**, 550 (2014).
49. Shah, V. P. et al. Bioanalytical method validation—a revisit with a decade of progress. *Pharm. Res.* **17**, 1551–1557 (2000).
50. Agren, R. et al. Reconstruction of genome-scale active metabolic networks for 69 human cell types and 16 cancer types using INIT. *PLoS Comput. Biol.* **8**, e1002518 (2012).
51. cobratoolbox. *GitHub* <https://github.com/opencobra/cobratoolbox/tree/master/src/dataIntegration/XomicsToModel> (2025).
52. COBRA.tutorials. *GitHub* https://opencobra.github.io/cobratoolbox/stable/tutorials/tutorial_XomicsToModel.html (2025).
53. Agren, R. et al. Identification of anticancer drugs for hepatocellular carcinoma through personalized genome-scale metabolic modeling. *Mol. Syst. Biol.* **10**, 721 (2014).
54. Gudmundsson, S. & Thiele, I. Computationally efficient flux variability analysis. *BMC Bioinform.* **11**, 489 (2010).
55. Noronha, A. et al. The virtual metabolic human database: integrating human and gut microbiome metabolism with nutrition and disease. *Nucleic Acids Res.* **47**, D614–D624 (2019).

Acknowledgements

We thank H. Leegwater for her helpful feedback while revising the manuscript. We acknowledge funding from the European Union's Horizon 2020 Programme (668738), Horizon Europe Framework Programme (101080997), European Research Council (101125633) and Dutch Research Council NWO project 184.034.019.

Author contributions

G.P.: conceptualization, methodology, software, tutorial and writing—original draft; A.W.: methodology, software, data curation and writing—review and editing; X.L.: methodology and software; I.T.: conceptualization, methodology, software and funding acquisition; T.H.: resources, supervision and funding acquisition; and R.M.T.F.: conceptualization, methodology, software, writing—original draft, writing—review and editing, supervision and funding acquisition.

Competing interests

The authors declare no competing interests.

Additional information

Supplementary information The online version contains supplementary material available at <https://doi.org/10.1038/s41596-025-01288-9>.

Correspondence and requests for materials should be addressed to Thomas Hankemeier or Ronan M. T. Fleming.

Peer review information *Nature Protocols* thanks Miguel Rocha and the other, anonymous reviewer(s) for their contribution to the peer review of this work.

Reprints and permissions information is available at www.nature.com/reprints.

Publisher's note Springer Nature remains neutral with regard to jurisdictional claims in published maps and institutional affiliations.

Springer Nature or its licensor (e.g. a society or other partner) holds exclusive rights to this article under a publishing agreement with the author(s) or other rightsholder(s); author self-archiving of the accepted manuscript version of this article is solely governed by the terms of such publishing agreement and applicable law.

© Springer Nature Limited 2025