Scalable metabolic pathway analysis

Ove Øyås[1,2] and Jörg Stelling[1,*]

[1] Department of Biosystems Science and Engineering and SIB Swiss Institute of
Bioinformatics, ETH Zurich, Basel, Switzerland
[2] Present address: Department of Animal and Aquacultural Sciences, Norwegian University of
Life Sciences, Ås, Norway
[*] Correspondence to: joerg.stelling@bsse.ethz.ch

The scope of application of genome-scale constraint-based models (CBMs) of metabolic networks rapidly expands toward multicellular systems. However, comprehensive analysis of CBMs through metabolic pathway analysis remains a major computational challenge because pathway numbers grow combinatorially with model sizes. Here, we define the minimal pathways (MPs) of a metabolic (sub)network as a subset of its elementary flux vectors. We enumerate or sample them efficiently using iterative minimization and a simple graph representation of MPs. These methods outperform the state of the art and they allow scalable pathway analysis for microbial and mammalian CBMs. Sampling random MPs from *Escherichia coli*'s central carbon metabolism in the context of a genome-scale CBM improves predictions of gene importance, and enumerating all minimal exchanges in a host-microbe model of the human gut predicts exchanges of metabolites associated with host-microbiota homeostasis and human health. MPs thereby open up new possibilities for the detailed analysis of large-scale metabolic networks.

Introduction

Constraint-based models (CBMs) can represent the complete genome-scale metabolic network of an organism and allow for the analysis of steady-state fluxes (reaction rates) based only on stoichiometry and other physico-chemical constraints. Large collections of curated models[1,2] and tools for automated model generation[3] make CBMs available for a wide range of organisms. Their scope of application is also expanding from unicellular to multicellular systems such as microbial communities or human tissues, increasing the need for analysis methods that scale to large metabolic networks[4].

The common analysis of CBMs assumes that metabolism is essentially at steady state on the time scale of other cellular processes. This produces a linear system of m metabolite mass balances in which the

variables are the n fluxes $\mathbf{r} \in \mathbb{R}^n$ constrained by the stoichiometric matrix $\mathbf{N} \in \mathbb{R}^{m \times n}$,

$$\mathbf{Nr} = 0, \tag{1}$$

and any other linear constraints, e.g., upper and lower flux bounds:

$$\mathbf{r}^{\min} \leq \mathbf{r} \leq \mathbf{r}^{\max}. \tag{2}$$

Solutions that satisfy these constraints are feasible combinations of steady-state fluxes known as flux distributions[5].

The most elementary method for analyzing flux distributions is flux balance analysis (FBA), in which an assumed objective is maximized or minimized subject to constraints (1) and (2) in a linear program (LP):

$$
\begin{aligned}
\text{max/min} \quad & \mathbf{w}^{\mathrm{T}}\mathbf{r} \\
\text{s. t.} \quad & \mathbf{Nr} = 0 \\
& \mathbf{r}^{\min} \leq \mathbf{r} \leq \mathbf{r}^{\max},
\end{aligned}
\tag{3}
$$

where $\mathbf{w}$ is the vector of weights for the fluxes. Maximal growth rate, as represented by the flux of a biomass reaction, is usually assumed as the objective for microbes[6]. FBA scales to very large networks thanks to efficient LP solvers, but it only returns a single optimal solution that is not necessarily unique[7].

In contrast to FBA and other biased methods, unbiased approaches can characterize the whole solution space. Random flux sampling can be used to obtain probability distributions over fluxes, but uniform sampling in high-dimensional spaces is challenging[8]. Pathway analysis is more comprehensive, defining all possible solutions in terms of sets of pathways such as elementary flux modes (EFMs)[9] or elementary flux vectors (EFVs)[10,11], as illustrated in **Fig. 1**. However, the number of pathways grows combinatorially with network size[12,13]. Hence, both approaches pose formidable challenges for large-scale analysis of metabolic networks.

Many conceptual and computational advances have been made to address these challenges. State-of-the-art algorithms can enumerate EFMs in networks with hundreds of reactions[14], and modularization enables enumeration of billions of EFMs[15]. To circumvent combinatorial explosion and to obtain sets of pathways amenable to analysis, one can limit enumeration to a subset of pathways[16]. Examples of such subsets are the shortest EFMs[17], random EFMs[18], and EFMs that are consistent with thermodynamics[19], regulation[20], or metabolomics data[21]. Elementary flux patterns (EFPs), defined as the parts of EFMs in a network that pass through a subnetwork (**Fig. 1**), have also been introduced[22]. However, while these approaches reduce the number of pathways, they do not scale to large networks.

EFMs are also limited because they are only defined for homogeneous constraints such as (1) or reaction reversibilities. Elementary flux vectors (EFVs) generalize EFMs by allowing inhomogeneous constraints such as (2) or any other linear constraint, but the challenges of enumeration remain[10,11]. Alternative

Figure 1. Example network with comparison of EFMs, EFVs, EFPs, and MPs. Two metabolites, A and B, participate in four reactions, r_1–r_4. The boundary reactions r_1 and r_4 exchange A and B with the environment and the internal reactions r_2 and r_3 define a subnetwork. All reactions except r_3 are irreversible, and we apply the maximum bound $\pm 1{,}000$ to all fluxes except r_4, which is fixed to 1. The pathways defined by the EFMs, EFVs, EFPs, and MPs of the (sub)network are shown below. Reactions that are not active in a pathway are not shown and active reactions that are not in the subnetwork are shown in gray for EFPs and MPs. EFMs and EFVs can be used to find all the MCSs of the network, and EFPs and MPs can be used to find all MCSs of the subnetwork, in this case $\{r_2, r_3\}$. The pathways span different subnetwork flux spaces: EFMs and EFPs (blue) do not take bounds into account and their flux space is unbounded in the direction of positive r_2 and r_3, while EFVs (dashed orange line) and MPs (red line) account for bounds and span bounded subspaces of the EFM/EFP space. There are fewer MPs than other pathways and the MPs are the only pathways that do not include loops ($r_3 < 0$).

approaches aim to find minimal sets of reactions[23–25] or precursor metabolites[26,27] under inhomogeneous constraints. However, they all rely on direct minimization of the same mixed-integer linear program (MILP), which prevents scaling to large networks. The same MILP is also used to compute minimal cut sets (MCSs), minimal sets of reactions that, when inactivated, block specific network functions[28].

Here, we introduce minimal pathways (MPs) as a concept well-suited for scalable pathway analysis (**Fig. 1**), and we present new methods for finding, sampling, and enumerating them in large metabolic (sub)networks. Our enumeration algorithm builds on recent work showing that the MILP can be split into an LP and a binary integer program (BIP), which can be alternated to compute EFMs and MCSs sequentially[29]. Our methods outperform the MILP approach when applied to a range of microbial and mammalian genome-scale CBMs, and we demonstrate applications to the analysis of central carbon metabolism in a genome-scale context and of host-microbe interactions in the human gut.

Results

Minimal pathways and detection methods

We define an MP as a support-minimal subset of reactions from a metabolic subnetwork, which in turn is a subset of reactions from a full metabolic network. Hence, all reactions in an MP need to be active (have non-zero flux) when no other reactions in the subnetwork are active in order to satisfy all constraints on the network as a whole. Combinations of fluxes in the subnetwork that satisfy all constraints on the full network lie in the projection of the solution space of the full network onto the subspace defined by the subnetwork (**Fig. 1**). The minimal set of vectors that generate this space without cancellations are its EFVs [10,11], and the subset of EFVs that are support-minimal is equivalent to the set of all MPs for a given subnetwork and network.

To find MPs efficiently, we developed a novel iterative minimization approach. It leverages support-minimality: to find an MP, we iteratively remove reactions from the subnetwork and minimize subnetwork flux with an LP until no further reactions can be removed without disrupting the functional requirements of the network. A randomized version allows random sampling of MPs. To enumerate MPs and MCSs, we alternate iterative minimization with computation of cut sets in a separate BIP, as recently suggested for EFMs [29]. We also realized that maximal cliques in a simple graph representation of known MPs can predict subnetworks likely to contain unknown MPs, thereby reducing the sizes of optimization problems during enumeration.

To illustrate the different approaches, we consider MP enumeration and sampling in a random subnetwork of 562 reactions in the genome-scale *E. coli* model iML1515 [30] with 2,868 irreversible reactions. Here, graph-based ennumeration found all 27,648 MPs within 45 minutes (**Fig. 2A**), but after one hour runtime, the other methods had found an order of magnitude fewer MPs (**Fig. 2B**). The graph-based approach also found all MCSs, primarily at iterations where the full subnetwork was used rather than a subset defined by a clique from the graph (**Fig. 2A,C**). Notably, iterative minimization without graph found all MCSs very quickly, while the graph-based approach found them more gradually (**Fig. 2C**). In more detail, direct minimization only identified MPs of the minimal size 14, iterative minimization without graph and randomization found MPs in arbitrary but clearly not random order, and randomized iterative minimization found MPs in apparently random order (**Fig. 2D**). The order of MP sizes for graph-based enumeration was more structured than for the other iterative approaches; there was a tendency to first find small MPs and to find many MPs of similar size after one another (**Fig. 2D**). The effective subnetwork size in graph-based enumeration varied between the full subnetwork when no new cliques were available from the graph and much smaller sizes approximately equal to the sizes of MPs when cliques were available (**Fig. 2E**). Specifically, the cliques from the graph corresponded exactly to MPs in 70% of the cases, and deviated only by one reaction from the MP otherwise. For the other methods, subnetwork size was constant and equal to the size of the full subnetwork.

Figure 2. Example enumeration with comparison of methods. (**A**) Complete enumeration of MPs and MCSs using iterative minimization with graph in a random subnetwork of size $|\mathcal{S}_0| = 562$ within the genome-scale *E. coli* model iML1515 with 2,868 irreversible reactions. The number of MPs $|\Pi|$ and MCSs $|\Gamma|$ found is shown along with the MP size $|\mathcal{P}|$ and effective subnetwork size $|\mathcal{S}|$ as a function of running time Δt. Values are normalized by the size of the complete set of MPs $|\Pi_c|$, the size of the complete set of MCSs $|\Gamma_c|$, or the initial subnetwork size $|\mathcal{S}_0|$. (**B-D**) Comparison to incomplete enumerations in the same subnetwork performed with direct minimization and iterative minimization with and without randomization, shown separately for $|\Pi|$ (**B**), $|\Gamma|$ (**C**), $|\mathcal{P}|$ (**D**), and $|\mathcal{S}|$ (**E**). End points of the completed enumeration are indicated by dots.

Benchmarking of methods

We benchmarked our algorithms using six models of microbial and mammalian cells, ranging in size from about 100 to more than 14,000 irreversible reactions (**Table S1**). Specifically, we sampled 100 subnetworks each of sizes ranging from 10 reactions to the full network (**Supplementary Data 1** and **Methods**). The running time needed to find the first MP was always greater and grew much faster with subnetwork size for direct than for iterative minimization, demonstrating that our method is faster and scales to larger networks than the state of the art (**Fig. 3A,B** and **Fig. S1**). Direct minimization failed to find a single MP in many of the largest subnetworks across all models except the smallest model of *E. coli* core metabolism (**Fig. S2A**). Randomized iterative minimization outperformed direct minimization across all subnetwork sizes in smaller models and the largest subnetworks in larger models, but it was slower for intermediate subnetwork sizes in larger models (**Fig. 3A,C**). In addition to subnetwork size, running time should depend on the sizes of the full networks in which the subnetworks were embedded.

Figure 3. Running times for finding the first MP. (**A**)-(**C**) Running time Δt_0 as a function of subnetwork size $|\mathcal{S}|$ for direct minimization (**A**), iterative minimization (**B**), and randomized iterative minimization (**C**). Lines indicate means with 95% confidence intervals from bootstrapping and colors indicate models. Enumeration was stopped after one hour (indicated by dashed line). (**D**) Running time Δt_0 as a function of subnetwork size $|\mathcal{S}|$, both normalized by network size $|\mathcal{N}|$, for direct (blue) and iterative (red) minimization. The double exponential model $\log y = a + b^x$ was fitted, yielding parameter estimates of $a = -4.97 \pm 0.02$ and $b = 7.32 \pm 0.13$ for direct minimization and $a = -5.50 \pm 0.01$ and $b = 2.33 \pm 0.02$ for iterative minimization without randomization.

Indeed, normalizing running time and subnetwork size by the size of the full network led to similar scaling across all models for direct minimization and iterative minimization without randomization according to a double exponential model (**Fig. 3D**). However, the running time of randomized iterative minimization scaled with absolute subnetwork size, following a monomial model (**Fig. 3C**).

Our graph-based approach completed more enumerations than any other method across all models, including all enumerations that were completed by any of the other methods (**Fig. S2B**). Reflecting the double exponential scaling of running time, the number of MPs found in enumerations that were completed within one hour grew superexponentially with subnetwork size (**Fig. S3A**). This was also the case for the number of MCSs obtained from enumerations using iterative minimization, but the set of all MCSs was generally smaller than the set of all MPs for subnetworks from all models except the (small) *E. coli* core model (**Fig. S3B**).

For enumerations that were not completed within one hour, the methods clearly differed in the number of MPs found. For direct minimization, the number of MPs decreased rapidly with subnetwork size for all models (**Fig. 4A** and **Fig. S3**). In contrast, iterative minimization without graph or randomization found approximately the same number of MPs (on the order of 10^3) in all subnetworks across all models (**Fig. 4B**). For randomized iterative minimization, MP numbers in incomplete enumerations increased with subnetwork size until a model-dependent critical size was reached, after which they decreased lin-

Figure 4. MPs found in incomplete enumerations. Number of MPs $|\Pi|$ found in one hour as a function of subnetwork size $|\mathcal{S}|$ for (**A**) direct minimization, (**B**) iterative minimization, (**C**) randomized iterative minimization, (**D**) and iterative minimization with graph. We excluded enumerations that were completed within one hour, except for randomized iterative minimization, where completion cannot be detected.

early with little variation between subnetworks and models (**Fig. 4C**). In most subnetworks, an order of magnitude more MPs were found by iterative minimization with graph than by any other method, but the number of MPs decreased more with subnetwork size when a graph was used (**Fig. 4B,D**). More MPs were found with randomized than with graph-based iterative minimization in some smaller subnetworks, but scaling was much worse with randomization (**Fig. 4C,D**).

Thus, iterative minimization was systematically faster than direct minimization and allowed scaling all the way to the full networks. Randomization allowed scalable random sampling and our graph-based approach accelerated enumeration.

Gene essentiality in central metabolism

To demonstrate the usefulness of MP sampling, we first analyzed gene essentiality, that is, whether a gene is required for the organism's survival and reproduction. This is not only a standard test scenario for CBMs, but gene essentiality is also increasingly recognized as quantitative and context-dependent[31]. Even for model organisms such as *E. coli* it still unclear what the essential genome is, with targeted knock-out libraries[32], randomized transposon mutagenesis[33], and most recently CRISPRi screens[34,35] yielding consensus as well as divergent gene sets.

We applied our randomized iterative minimization procedure to the central metabolism of *E. coli* as a subnetwork in the genome-scale model iJO1366[36]. This subnetwork of 153 irreversible reactions, mostly

Figure 5. Example application of MP sampling. (**A**) Pathway map of *E. coli* central carbon metabolism. Circles are metabolites and arrows are reactions. The color and width of arrows indicate the fraction of 100,000 randomly sampled MPs in which reactions participated. The network was used as a subnetwork within the genome-scale model iJO1366 with a growth requirement of 0.1 h^{-1}, a minimal glucose medium, and aerobic conditions. For reversible reactions, the fraction includes all MPs with the reaction regardless of flux direction. (**B**) Optical density (OD) of gene knockout strain at 24 h as a function of gene importance predicted by MP sampling (also indicated by color). A fitted line is shown with 95% confidence interval from bootstrapping. (**C**) Absolute correlation of gene importance predicted by MP sampling (red) and FBA (blue) with growth (OD of knockout strain at 24 h) and experimental essentiality score. Stars indicate significance ($p < 0.05$).

from glycolysis, the tricarboxylic acid (TCA) cycle, and the pentose phosphate pathway, is embedded in the full model of 3,219 irreversible reactions. Specifically, we sampled 100,000 MPs (in 10.3 h with 0.37 ± 0.18 s per iteration) for aerobic growth on a minimal glucose medium.

To reveal the relative importance of reactions and enzymes in central carbon metabolism, we calculated the fraction of MPs in which each reaction or enzyme participates (**Fig. 5A**, **Supplementary Data 2 and 3**, and **Methods**). No enzymes in the pentose phosphate pathway were essential, but five enzymes participated in more than 80% of all MPs (transaldolase, transketolases, ribulose 5-phosphate 3-epimerase, and ribose-5-phosphate isomerase). Also, four enzymes stood out in the TCA cycle (citrate synthase, aconitase, isocitrate dehydrogenase, and fumarase), appearing to be much more important for aerobic growth on glucose than, for example, 2-oxogluterate dehydrogenase and succinyl-CoA synthetase. As expected, isozymes such as the malate dehydrogenases in the TCA cycle were of very similar and intermediate importance. These patterns are consistent with the alternative operation of the full TCA cycle and the PEP-glyoxylate cycle for complete glucose oxidization[37].

To evaluate our predictions quantitatively, we next mapped samples to enzyme-encoding genes and compared the relative importance of the genes to available growth and gene essentiality data from the Keio collection[32]. Gene importance in terms of growth reduction upon gene deletions predicted by MP sampling of central carbon metabolism correlated as well with growth of knock-out strains as FBA predictions for the full model (**Fig. 5B** and **Fig. S5**; Pearson's $r = -0.46$, $p = 2 \cdot 10^{-10}$ for MPs). Also, both approaches yielded identical results for the binary classification of essential and non-essential reactions based on knock-out strains. In a previous elementary mode analysis with reduced models, however, small deviations in predictions were observed[38]. Finally, we analyzed the methods' power of predicting Keio essentiality scores; they result from aggregating targeted knock-out and transposon mutagenesis data sets and range from -4 (indicating non-essentiality) to $+3$ (indicating essentiality). MP sampling clearly outperformed growth predictions from FBA in recapitulating these measures of experimental gene essentiality (**Fig. 5C** and **Fig. S5**), indicating that MP analysis may better reflect context-dependency in the data than point estimates obtained by FBA.

Metabolite exchanges in the human gut

To show that our graph-based approach enables complete enumeration of MPs in large models of multicellular systems, we applied it to a host-microbe model of the human gut. We know little about how different members of the gut microbiota interact with each other and with the host, in particular with respect to the metabolite exchanges involved[39]. However, experimental studies have shown that metabolic pathway- rather than species-associations are predictive of host-microbe metabolic interactions[40,41], making pathway analysis a suitable approach. Yet, whereas computational tools for the analysis of larger microbial communities have been developed (mostly as variants of FBA), applications are missing so far in general[39], and no comprehensive pathway analysis exists specifically.

Figure 6. Example application of MP enumeration. (**A**) Complete enumeration of MPs and MCSs using iterative minimization with graph in a subnetwork of 3,209 metabolite exchanges within a host-microbe model with 33,291 irreversible reactions (symbols used as in **Fig. 2A**). (**B**) Fraction of MPs with uptake (red) and secretion (blue) for each metabolite (columns) and organism (rows).

In our host-microbe model (see **Methods** for details), a human cell interacts with six commensal microbial species. We required ATP production for the human cell and growth for the microbial cells in anaerobic conditions and a complex metabolic environment. We used our recent approach to constrain the model to explain as much of the functional requirements as possible in terms of metabolite exchanges[42], and enumerated all minimal sets of metabolite exchanges between microbes and from microbes to the human cell using iterative minimization with graph.

In this subnetwork of 3,209 metabolite exchange reactions within the full network of 33,291 irreversible reactions, we found all 240 MPs and 63 MCSs in 0.52 h with 7.6 ± 3.3 s per iteration (**Fig. 6A**). Because we constrained the problem to functionally required exchanges for given growth conditions, the MP set was small, indicating a focus on important exchanges (relaxations yield larger sets of MPs, but this was not the focus of our study). More specifically, the MPs were different combinations of 17 uptakes and 18 secretions, corresponding to 28 unique pairwise exchanges of ten different metabolites: amino acids, fermentation products, TCA intermediates, and the B vitamin niacin (**Fig. 6B**). The human cell was predicted to consume microbial formate, glycine, glycerol, malate, and ornithine, and the microbes exchanged all metabolites except formate and glycerol.

Importantly, the metabolites predicted to be exchanged have been observed in the gut and implicated in the maintenance of host-microbiota homeostasis as well as human health and disease. For example, a possible alanine exchange predicted by the same microbial model collection used here was confirmed *in vitro*[43], increased isoleucine levels have been associated with improved growth in humans[44], and consumption of glycine by the microbiota has been shown to reduce glycine levels in the host, which has been associated with non-alcoholic fatty liver disease (NAFLD), obesity, and type 2 diabetes[45]. The gut microbiota also contributes more directly to NAFLD by producing ethanol, which suggests that consumption of ethanol through cross-feeding could be important in healthy gut microbiotas[46]. Glycerol degradation pathways are enriched in gut microbiotas of people with functional constipation[47], elevated levels and cross-feeding of microbially derived formate have been associated with inflammation[48], and levels of succinate as well as microbes producing or consuming it have been associated with obesity and

related diseases[49]. Exchange of niacin (B3) and its importance for microbiota stability are supported by *in silico*, *in vitro*, and *in vivo* evidence[40,50,51], and microbial ornithine production promotes healthy gut mucosa through crosstalk with the host[52]. Thus, the predicted exchanges appear plausible, and they amount to hypotheses on specific producers and consumers, which are hard to identify experimentally.

Discussion

Our definition of MPs in terms of EFVs makes it straight-forward to reconcile MPs with available theory and tools, while providing several advantages over existing pathway definitions. First, as EFVs, MPs account for all linear constraints on the network as a whole, including inhomogeneous constraints neglected by EFMs. However, because MPs are a subset of the EFVs, there are fewer MPs than EFVs to enumerate. Second, all reactions in an MP contribute to fulfilling functional requirements as defined by constraints on the full network, which is not the case for EFMs, EFPs, or all EFVs. While this implies that MPs do not account for all possible flux distributions of the subnetwork, they cover the minimal flux distributions that support function, which have been shown to be consistent with omics data[53,54]. Also, MPs avoid the pervasive problem of thermodynamically infeasible loops[55], which cannot contribute to well-defined functional requirements. Third, like other pathway definitions, MPs can be used to find all MCSs in the subnetwork that knock out the function of the full network. Finally, like EFPs, MPs are defined for subnetworks as well as full networks, which allows one to adjust the subnetwork size for complete pathway enumeration in practice.

We argue that our iterative minimization approach based on LPs can replace the computationally hard direct minimization of a single MILP in the state of the art for pathway enumeration and sampling[18,22,25]. Running time as a function of subnetwork size, normalized by the size of the full network, was described well by a double exponential model for both approaches. However, our iterative approach had a constant that was an order of magnitude smaller than the direct approach, and it scaled with a much smaller base for exponentiation, allowing MPs to be found in networks of all sizes. More specifically, our graph-based approach completed more enumerations within one hour than other methods, and it yielded an order of magnitude more MPs in most incomplete enumerations. However, the number of MPs found decreased slightly with subnetwork size, likely due to the identification of maximal cliques, which is computationally hard and suffers from the increasing size of the graph with the number of MPs[56]. Here, we detected all maximal cliques each time the graph was updated (when the full subnetwork was used to find an MP), but future implementations could benefit from directly finding cliques with edges between reactions in the latest MP.

For randomized iterative minimization, absolute running time scaled with absolute subnetwork size, reflecting that it needs to solve a number of LP problems proportional to subnetwork size. This would also explain why direct minimization outperformed randomized iterative minimization in subnetworks large in absolute terms, but small relative to the full network. Note that no complete enumerations were

registered by the randomized algorithm because we did not check for completeness; this would require integer variables and constraints at significant performance cost. We emphasize that the main bottleneck for enumeration is the BIP used to find cut sets, which slows down each iteration, especially as the number of known MPs and MCSs increases. Thus, a faster procedure for finding new cut sets would be a major contribution toward further scaling. As a general guideline for applying our methods to (sub)networks of interest, we therefore suggest that graph-based enumeration is used first to check if MPs and MCSs are enumerable in reasonable time. If not, the subnetwork size can be reduced, additional constraints can be applied, or MPs can be sampled randomly rather than enumerated. We also recommend network compression, as commonly applied before EFM computation[14], to reduce the initial problem sizes as much as possible.

We demonstrated that random sampling and enumeration of MPs can yield biological insights through two example applications. First, random sampling of MPs in *E. coli* central carbon metabolism within a genome-scale network revealed the relative importance of reactions and enzyme-encoding genes in better agreement with experimental gene essentiality data than FBA. This approach is straight-forward and could be valuable for many applications, for example for the selection of knock-outs unlikely to disrupt growth or product formation in metabolic engineering strategies. Second, we enumerated all minimal sets of metabolite exchanges in a host-microbe model of a human cell interacting with six gut microbes. This showed that graph-based enumeration can be used to analyze some of the largest models considered to date, and our predictions included metabolite exchanges with experimental evidence and possible roles in gut homeostasis and human health. We note that future studies of gut microbiotas could use relaxed problem formulations to generate more MPs for potential metabolite exchanges—here, we aimed to identify minimal requirements for validation against the experimental literature. Overall, we argue that MPs open up new possibilities for detailed analysis of large-scale metabolic networks, and we expect them to produce new biological insights into uni- and multicellular metabolic systems.

Methods

Minimal pathways and cut sets

We start from a CBM with a flux vector $\mathbf{r}$ constrained by (1) and (2). We define the metabolic network $\mathcal{N} = \{1, 2, \ldots, n\}$ as the set of reaction indices in $\mathbf{r}$, and we define a subnetwork $\mathcal{S}$ and a pathway $\mathcal{P}$ such that $\mathcal{P} \subseteq \mathcal{S} \subseteq \mathcal{N}$. We further require $\mathcal{P}$ to be a minimal pathway (MP), here defined as a minimal set of reactions in $\mathcal{S}$ that need to be active (have non-zero flux) in order to satisfy (1), (2), and any other linear constraints on $\mathcal{N}$. In other words, $\mathcal{P}$ is *support-minimal*: if the only active reactions in $\mathcal{S}$ are $\mathcal{P}$, constraints on $\mathcal{N}$ will be violated if any reaction in $\mathcal{P}$ is deactivated. We denote the set of known MPs as Π.

We define a cut set $\mathcal{C} \subseteq \bigcup \Pi$ as a set of reactions that intersects every $\mathcal{P} \in \Pi$. This means that deactivating all reactions in $\mathcal{C}$ will deactivate all known MPs. We also require $\mathcal{C}$ to be support-minimal such that removing any reaction from $\mathcal{C}$ will leave an empty intersection with at least one $\mathcal{P} \in \Pi$. A minimal cut set (MCS) deactivates all MPs in $\mathcal{S}$ and thereby disables flux in $\mathcal{N}$ as a whole. We denote the set of known MCSs as Γ.

Model preprocessing

We use flux variability analysis (FVA;[7]) to determine tight bounds ($r_i^{\min}$, $r_i^{\max}$) on feasible fluxes for all $i \in \mathcal{N}$. Reactions in $\mathcal{S}$ are then made irreversible with positive upper bounds by replacing reversible reactions by two irreversible ones and reversing irreversible reactions with negative lower bounds. FVA also reveals blocked reactions,

$$\mathcal{B} = \{i \in \mathcal{N} \mid r_i^{\max} = 0\}, \tag{4}$$

which always have zero flux and therefore are not part of any MP or MCS, and essential reactions,

$$\mathcal{E} = \{i \in \mathcal{N} \mid r_i^{\min} > 0\}, \tag{5}$$

which cannot have zero flux and are MCSs that participate in all MPs. Hence, only the subnetwork $\mathcal{S} \setminus (\mathcal{B} \cup \mathcal{E})$ can form MPs and MCSs along with $\mathcal{E}$. For simplicity, we will keep denoting this subnetwork as $\mathcal{S}$.

Finding a new pathway with direct minimization

For each $i \in \mathcal{N}$, we use a binary variable $b_i \in \{0, 1\}$ to indicate if r_i is non-zero, as ensured by an inequality constraint:

$$r_i - r_i^{\max} b_i \leq 0. \tag{6}$$

Another inequality constraint stops each $\mathcal{P} \in \Pi$ from being found again:

$$\sum_{i \in \mathcal{P}} b_i \leq |\mathcal{P}| - 1 \tag{7}$$

We can combine constraints (2), (6), and (7):

$$\mathbf{A} \begin{bmatrix} \mathbf{r} \\ \mathbf{b} \end{bmatrix} \leq \mathbf{c}, \tag{8}$$

where the matrix $\mathbf{A}$ contains the coefficients, $\mathbf{b}$ is the vector of indicator variables, and $\mathbf{c}$ is the vector of constant terms. A new MP is then found by solving the MILP

$$\begin{aligned} \min \quad & \sum_{i \in \mathcal{S}} b_i \\ \text{s. t.} \quad & \mathbf{Nr} = 0 \\ & \mathbf{A} \begin{bmatrix} \mathbf{r} \\ \mathbf{b} \end{bmatrix} \leq \mathbf{c}. \end{aligned} \tag{9}$$

It minimizes the number of active reactions in $\mathcal{S}$ by minimizing the number of non-zero indicator variables. In an optimal solution, the active reactions in $\mathcal{S}$ define an MP that has the smallest possible number of active reactions among all MPs that have not yet been found.

Finding a new cut set

We find a new cut set that deactivates all known MPs as described for EFMs by Song et al.[29]. Each reaction in at least one known MP, $i \in \bigcup \Pi$, is represented by a binary variable $b_i \in \{0, 1\}$, and we use a constraint for each $\mathcal{P} \in \Pi$ to ensure that it is covered by the cut set:

$$\sum_{i \in \mathcal{P}} b_i \geq 1. \tag{10}$$

We also use a constraint for each $\mathcal{C} \in \Gamma$ to exclude known MCSs:

$$\sum_{i \in \mathcal{C}} b_i \leq |\mathcal{C}| - 1. \tag{11}$$

By reformulating constraints (10) and (11) as in (8), a new cut set results from solving the binary integer program (BIP)

$$\begin{aligned} \min \quad & \sum_{i \in \bigcup \Pi} b_i \\ \text{s. t.} \quad & \mathbf{Ab} \leq \mathbf{c}. \end{aligned} \tag{12}$$

Specifically, a new cut set is given by the non-zero binary variables in the optimal solution; it may or may not be an MCS.

To find a new MP with iterative minimization, it is necessary to find a *viable* cut set that intersects all $\mathcal{P} \in \Pi$ but does not intersect at least one unknown MP, i.e., a cut set that is not an MCS. To achieve this, **Algorithm S1** repeatedly solves and constrains the BIP and checks viability with the LP.

Finding a new pathway with iterative minimization

Assuming that a cut set that disables known MPs has been applied, our iterative method (**Algorithm S2**) does not require integer variables or additional constraints to find a new MP. At each iteration, a reaction $i \in \mathcal{S}$ is deactivated before minimizing the sum of fluxes in $\mathcal{S}$ in the LP

$$
\begin{aligned}
\min \quad & \sum_{i \in \mathcal{S}} r_i \\
\text{s. t.} \quad & \mathbf{Nr} = 0 \\
& \mathbf{r}^{\min} \leq \mathbf{r} \leq \mathbf{r}^{\max}.
\end{aligned}
\tag{13}
$$

If an optimal solution is found, reactions in $\mathcal{S}$ with zero flux in the solution can be deactivated as well because they are not needed in $\mathcal{P}$. If no optimal solution is found, i must be part of $\mathcal{P}$ and is reactivated. We repeat this until all $i \in \mathcal{S}$ have been tested, and the remaining active reactions in $\mathcal{S}$ constitute an MP that is support-minimal but does not necessarily have the smallest possible number of reactions.

We can also randomize the order of MPs by deactivating a random $i \in \mathcal{S}$ at each iteration. In this case, we cannot deactivate all reactions in $\mathcal{S}$ with zero flux in an optimal solution because randomness would not be preserved. However, random reactions can be drawn and deactivated until a reaction with non-zero flux in the optimal solution is drawn.

Simple enumeration and sampling of pathways

Direct minimization can be used to enumerate MPs in order of increasing size by repeating the following steps until the MILP becomes infeasible:

1. Find a new MP by solving the MILP.

2. Add constraint (7) to the MILP to avoid finding the MP again.

Iterative minimization can be combined with the BIP for finding new cut sets to enumerate MPs based on the procedure by Song et al.[29]. The following steps are repeated until the BIP becomes infeasible:

1. Find and apply a new viable cut set using **Algorithm S1**.

2. Find a new MP using **Algorithm S2**.

3. Add constraint (10) to the BIP to avoid finding the cut set again.

This produces all MPs (not in order of increasing size) as well as all MCSs.

If too many MPs exist for complete enumeration to be feasible, MPs can be sampled randomly by repeatedly finding random MPs without finding and applying cut sets until the number of MPs equals the desired sample size. This avoids the significant overhead of solving BIPs but risks finding the same MP multiple times if the number of MPs is not sufficiently large.

Graph-based enumeration of pathways

An undirected graph can be used to predict new MPs from known MPs. We define this graph as the union of reaction pairs from all $\mathcal{P} \in \Pi$, $\mathcal{G} = \bigcup_{\mathcal{P} \in \Pi} \binom{\mathcal{P}}{2}$, where reactions are nodes and reaction pairs are edges connecting two nodes. Reactions are connected if they occur together in at least one known MP, meaning that MPs are *cliques* (complete subgraphs) in $\mathcal{G}$. We exploit that edges defined by known MPs can combine with each other to form additional cliques that may contain previously unknown MPs.

Each time we find a new MP during enumeration, we update the graph by adding any missing edges between reactions in the new MP. We use a modified Bron-Kerbosch algorithm[56] to find *maximal cliques*, i.e., fully connected subgraphs where no other nodes in the graph are connected to all nodes in the subgraph. The MPs in each new maximal clique are enumerated separately using iterative minimization with the clique as subnetwork in place of $\mathcal{S}$. This reduces the sizes of the optimization problems and thereby accelerates enumeration.

Graph-based prediction of new MPs can be combined with direct or iterative minimization (or any other procedure that returns new MPs); here, we implemented iterative minimization as described in **Algorithm S3**. However, graph-based prediction is biased by the known MPs and therefore incompatible with random sampling of MPs.

Benchmarking

We obtained six CBMs (**Table S1**) from the BiGG database[2], comprising models of *Escherichia coli*[30,57], *Helicobacter pylori*[58], *Saccharomyces cerevisiae*[59], *Cricetulus griseus*[60], and *Homo sapiens*[61]. We removed any non-growth-associated maintenance reactions and all biomass reactions except the one with the largest number of metabolites, set the default bounds of fluxes to $\pm$ 1,000 mmol gDW^{-1} h^{-1}, set a minimal growth rate of 0.1 h^{-1}, and allowed uptake and secretion of all environmental metabolites to simulate a complex medium. We preprocessed each model as described above.

For each model, we sampled random subnetworks ranging from ten reactions to the full network. We chose logarithmically spaced subnetwork sizes with four sizes sampled for each order of magnitude between 10 and 10^4. We sampled 100 random subnetworks of each size and enumerated MPs using direct and iterative minimization (with and without graph or randomization). Incomplete enumerations were stopped after one hour.

We used bootstrapping to compute 95% confidence intervals for estimated means. The data were resampled 1,000 times with replacement, the mean was reestimated for each sample, and a 95% confidence interval was defined from the 2.5th and 97.5th percentiles of these estimates.

Application to central metabolism

We obtained a pathway map of *E. coli* central carbon metabolism from Escher[62] and used the reactions in this map as a subnetwork within the genome-scale model iJO1366[36]. Using randomized iterative minimization, we sampled 100,000 MPs in the subnetwork with a growth requirement of 0.1 h^{-1}, the default minimal glucose medium, and aerobic conditions. We defined reaction (gene) importance as the fraction of sampled MPs requiring each reaction (enzyme-encoding gene). To map samples from reactions to genes, we computed the contribution of each gene to each reaction by dividing the number of enzymes requiring the gene by the total number of enzymes catalyzing the reaction. To predict gene importance with FBA, we deleted each gene from the model, maximized growth rate, and computed the fractional deviation from the maximal growth rate of the wild type. We obtained growth and gene essentiality data from the Keio collection[32]. Bootstrapping was performed as for benchmarking.

Application to host-microbe exchanges

We built a host-microbe model in which genome-scale CBMs of a human cell and of six commensal gut microbes (*Bacteroides thetaiotaomicron*, *E. coli*, *Faecalibacterium prausnitzii*, *Lactobacillus plantarum*, *Lactococcus lactis*, and *Streptococcus thermophilus*; **Table S2**) can exchange metabolites through the gut lumen. We based the model on previous work[63] but used a more recent human model[61] and collection of gut microbe models[43], all obtained from Virtual Metabolic Human[64]. We set a non-growth-associated ATP requirement of 0.1 mmol gDW^{-1} h^{-1} for the human cell and a growth requirement of 0.1 h^{-1} for the microbes. To simulate the anaerobic and potentially complex metabolic environment of the gut, we allowed uptake and secretion of all environmental metabolites for all cells, turned off oxygen exchange for the microbes, and added reactions representing all possible metabolite exchanges between microbes and from microbes to the human cell (excluding ATP, sugars, and inorganic ions and compounds). Building on our recent approach[42], we also applied two FBA-steps, minimizing (1) microbial intracellular flux[53] and (2) uptake flux not explained by metabolite exchange across all cells:

$$
\begin{aligned}
&\min \quad \|\mathbf{w}^{\mathrm{T}}\mathbf{r}\|_1 \\
&\text{s. t.} \quad \mathbf{N}\mathbf{r} = 0 \\
&\phantom{\text{s. t.} \quad} \mathbf{r}^{\min} \leq \mathbf{r} \leq \mathbf{r}^{\max},
\end{aligned}
\tag{14}
$$

where $\|\mathbf{w}^{\mathrm{T}}\mathbf{r}\|_1$ is the weighted sum of absolute fluxes to be minimized. After the first step, we constrained each intracellular microbial flux to its value in the minimum obtained. After the second step, we constrained the absolute sum of exchange fluxes not explained by metabolite exchange to the minimum

obtained, $\|\mathbf{w}^T\mathbf{r}\|_1 = \|\mathbf{w}^T\mathbf{r}\|_1^{\min}$. This approach aimed to explain as much of the functional requirements as possible in terms of metabolite cross-feeding with biologically reasonable microbial flux distributions. We enumerated minimal sets of metabolite exchanges using iterative minimization with graph.

Software

We implemented the algorithms in Python using COBRApy[65] and the Gurobi Optimizer (Gurobi Optimization, LLC, Beaverton, OR, USA).

Data and code availability

The data and code needed to reproduce our results are available at https://gitlab.com/csb.ethz/mptool.

Acknowledgements

This work was supported by the SystemsX.ch project GutX, evaluated by the Swiss National Science Foundation, and the Swiss National Science Foundation Sinergia project #177164. We thank Mattia Gollub, Axel Theorell, and the DigiSal team at the Norwegian University of Life Sciences for helpful comments on the manuscript.

Author contributions

J.S. and O.Ø. designed the study and wrote the manuscript. O.Ø. developed the algorithms and performed the analyses.

Conflict of interest

The authors declare no conflict of interest.

References

1. M. Ganter, T. Bernard, S. Moretti, J. Stelling, and M. Pagni. MetaNetX.org: a website and repository for accessing, analysing and manipulating metabolic networks. *Bioinformatics*, 29(6), 2013.

2. Z. A. King, J. Lu, A. Dräger, P. Miller, S. Federowicz, J. A. Lerman, A. Ebrahim, B. O. Palsson, and N. E. Lewis. BiGG Models: A platform for integrating, standardizing and sharing genome-scale models. *Nucleic Acids Research*, 44(D1), 2016.

3. S. N. Mendoza, B. G. Olivier, D. Molenaar, and B. Teusink. A systematic assessment of current genome-scale metabolic reconstruction tools. *Genome Biology*, 20(1), 2019.

4. O. Øyås and J. Stelling. Genome-scale metabolic networks in time and space. *Current Opinion in Systems Biology*, 8, 2018.

5. A. Varma and B. O. Palsson. Metabolic flux balancing: Basic concepts, scientific and practical use. *Bio/Technology*, 12(10), 1994.

6. A. M. Feist and B. O. Palsson. The biomass objective function. *Current Opinion in Microbiology*, 13(3), 2010.

7. R. Mahadevan and C. H. Schilling. The effects of alternate optimal solutions in constraint-based genome-scale metabolic models. *Metabolic Engineering*, 5(4), 2003.

8. J. Schellenberger and B. Palsson. Use of randomized sampling for analysis of metabolic networks. *Journal of Biological Chemistry*, 284(9), 2009.

9. S. Schuster and C. Hilgetag. On elementary flux modes in biochemical systems at steady state. *Journal of Biological Systems*, 2(2), 1994.

10. R. Urbanczik. Enumerating constrained elementary flux vectors of metabolic networks. *IET Systems Biology*, 1(5), 2007.

11. S. Klamt, G. Regensburger, M. P. Gerstl, C. Jungreuthmayer, S. Schuster, R. Mahadevan, J. Zanghellini, and S. Müller. From elementary flux modes to elementary flux vectors: Metabolic pathway analysis with arbitrary linear flux constraints. *PLOS Computational Biology*, 13(4), 2017.

12. S. Klamt and J. Stelling. Combinatorial complexity of pathway analysis in metabolic networks. *Molecular Biology Reports*, 29(1-2), 2002.

13. V. Acuña, F. Chierichetti, V. Lacroix, A. Marchetti-Spaccamela, M. F. Sagot, and L. Stougie. Modes and cuts in metabolic networks: Complexity and algorithms. *BioSystems*, 95(1), 2009.

14. M. Terzer and J. Stelling. Large-scale computation of elementary flux modes with bit pattern trees. *Bioinformatics*, 24(19), 2008.

15. K. A. Hunt, J. P. Folsom, R. L. Taffs, and R. P. Carlson. Complete enumeration of elementary flux modes through scalable demand-based subnetwork definition. *Bioinformatics*, 30(11), 2014.

16. J. Zanghellini, M. P. Gerstl, M. Hanscho, G. Nair, G. Regensburger, S. Müller, and C. Jungreuthmayer. Toward Genome-Scale Metabolic Pathway Analysis. In *Industrial Biotechnology*. 2016.

17. L. F. de Figueiredo, A. Podhorski, A. Rubio, C. Kaleta, J. E. Beasley, S. Schuster, and F. J. Planes. Computing the shortest elementary flux modes in genome-scale metabolic networks. *Bioinformatics*, 25(23), 2009.

18. D. Machado, Z. Soons, K. R. Patil, E. C. Ferreira, and I. Rocha. Random sampling of elementary flux modes in large-scale metabolic networks. *Bioinformatics*, 28(18), 2012.

19. S. J. Jol, A. Kümmel, M. Terzer, J. Stelling, and M. Heinemann. System-level insights into yeast metabolism by thermodynamic analysis of elementary flux modes. *PLOS Computational Biology*, 8(3), 2012.

20. C. Jungreuthmayer, D. E. Ruckerbauer, and J. Zanghellini. RegEfmtool: Speeding up elementary flux mode calculation using transcriptional regulatory rules in the form of three-state logic. *BioSystems*, 113(1), 2013.

21. M. P. Gerstl, D. E. Ruckerbauer, D. Mattanovich, C. Jungreuthmayer, and J. Zanghellini. Metabolomics integrated elementary flux mode analysis in large metabolic networks. *Scientific Reports*, 5(8930), 2015.

22. C. Kaleta, L. F. De Figueiredo, and S. Schuster. Can the whole be less than the sum of its parts? Pathway analysis in genome-scale metabolic networks using elementary flux patterns. *Genome Research*, 19(10), 2009.

23. A. P. Burgard, S. Vaidyaraman, and C. D. Maranas. Minimal reaction sets for Escherichia coli metabolism under different growth requirements and uptake environments. *Biotechnology Progress*, 17(5), 2001.

24. S. Jonnalagadda and R. Srinivasan. An efficient graph theory based method to identify every minimal reaction set in a metabolic network. *BMC Systems Biology*, 8(1), 2014.

25. J. Pey and F. J. Planes. Direct calculation of elementary flux modes satisfying several biological constraints in genome-scale metabolic networks. *Bioinformatics*, 30(15), 2014.

26. V. Acuña, P. V. Milreu, L. Cottret, A. Marchetti-Spaccamela, L. Stougie, and M. F. Sagot. Algorithms and complexity of enumerating minimal precursor sets in genome-wide metabolic networks. *Bioinformatics*, 28(19), 2012.

27. R. Andrade, M. Wannagat, C. C. Klein, V. Acuña, A. Marchetti-Spaccamela, P. V. Milreu, L. Stougie, and M. F. Sagot. Enumeration of minimal stoichiometric precursor sets in metabolic networks. *Algorithms for Molecular Biology*, 11(1), 2016.

28. A. von Kamp and S. Klamt. Enumeration of Smallest Intervention Strategies in Genome-Scale Metabolic Networks. *PLOS Computational Biology*, 10(1), 2014.

29. H. S. Song, N. Goldberg, A. Mahajan, and D. Ramkrishna. Sequential computation of elementary modes and minimal cut sets in genome-scale metabolic networks using alternate integer linear programming. *Bioinformatics*, 33(15), 2017.

30. J. M. Monk, C. J. Lloyd, E. Brunk, N. Mih, A. Sastry, Z. King, R. Takeuchi, W. Nomura, Z. Zhang, H. Mori, A. M. Feist, and B. O. Palsson. iML1515, a knowledgebase that computes Escherichia coli traits. *Nature Biotechnology*, 35(10), 2017.

31. G. Rancati, J. Moffat, A. Typas, and N. Pavelka. Emerging and evolving concepts in gene essentiality. *Nature Reviews Genetics*, 19(1), 2018.

32. T. Baba, T. Ara, M. Hasegawa, Y. Takai, Y. Okumura, M. Baba, K. A. Datsenko, M. Tomita, B. L. Wanner, and H. Mori. Construction of Escherichia coli K-12 in-frame, single-gene knockout mutants: The Keio collection. *Molecular Systems Biology*, 2(1), 2006.

33. E. C. Goodall, A. Robinson, I. G. Johnston, S. Jabbari, K. A. Turner, A. F. Cunningham, P. A. Lund, J. A. Cole, and I. R. Henderson. The essential genome of Escherichia coli K-12. *mBio*, 9(1), 2018.

34. T. Wang, C. Guan, J. Guo, B. Liu, Y. Wu, Z. Xie, C. Zhang, and X.-H. Xing. Pooled CRISPR interference screening enables genome-scale functional genomics study in bacteria with superior performance. *Nature Communications*, 9(1), 2018.

35. F. Rousset, L. Cui, E. Siouve, C. Becavin, F. Depardieu, and D. Bikard. Genome-wide CRISPR-dCas9 screens in E. coli identify essential genes and phage host factors. *PLoS Genetics*, 14(11), 2018.

36. J. D. Orth, T. M. Conrad, J. Na, J. A. Lerman, H. Nam, A. M. Feist, and B. Palsson. A comprehensive genome-scale reconstruction of Escherichia coli metabolism-2011. *Molecular Systems Biology*, 7 (535), 2011.

37. E. Fischer and U. Sauer. A novel metabolic cycle catalyzes glucose oxidation and anaplerosis in hungry Escherichia coli. *The Journal of Biological Chemistry*, 278(47), 2003.

38. O. Hädicke and S. Klamt. EColiCore2: a reference network model of the central metabolism of Escherichia coli and relationships to its genome-scale parent model. *Scientific Reports*, 7(39647), 2017.

39. M. Kumar, B. Ji, K. Zengler, and J. Nielsen. Modelling approaches for studying the microbiome. *Nature Microbiology*, 4(8), 2019.

40. A. Visconti, C. I. Le Roy, F. Rosa, N. Rossi, T. C. Martin, R. P. Mohney, W. Li, E. de Rinaldis, J. T. Bell, J. C. Venter, K. E. Nelson, T. D. Spector, and M. Falchi. Interplay between the human gut microbiome and host metabolism. *Nature Communications*, 10(1), 2019.

41. B. Yilmaz, P. Juillerat, O. Øyås, C. Ramon, F. D. Bravo, Y. Franc, N. Fournier, P. Michetti, C. Mueller, M. Geuking, V. E. H. Pittet, M. H. Maillard, G. Rogler, S. I. C. Investigators, R. Wiest, J. Stelling, and A. J. Macpherson. Microbial network disturbances in relapsing refractory Crohn's disease. *Nature Medicine*, 25(2), 2019.

42. O. Øyås, S. Borrell, A. Trauner, M. Zimmermann, J. Feldmann, T. Liphardt, S. Gagneux, J. Stelling, U. Sauer, and M. Zampieri. Model-based integration of genomics and metabolomics reveals SNP functionality in Mycobacterium tuberculosis. *Proceedings of the National Academy of Sciences*, 117 (15), 2020.

43. S. Magnúsdóttir, A. Heinken, L. Kutt, D. A. Ravcheev, E. Bauer, A. Noronha, K. Greenhalgh, C. Jäger, J. Baginska, P. Wilmes, R. M. Fleming, and I. Thiele. Generation of genome-scale metabolic reconstructions for 773 members of the human gut microbiota. *Nature Biotechnology*, 35(1), 2017.

44. M. Kumar, B. Ji, P. Babaei, P. Das, D. Lappa, G. Ramakrishnan, T. E. Fox, R. Haque, W. A. Petri, F. Bäckhed, and J. Nielsen. Gut microbiota dysbiosis is associated with malnutrition and reduced plasma amino acid levels: Lessons from genome-scale metabolic modeling. *Metabolic Engineering*, 49, 2018.

45. A. Mardinoglu, M. Bergentall, P. Ghaffari, E. Larsson, F. Backhed, S. Shoaie, J. Nielsen, and C. Zhang. The gut microbiota modulates host amino acid and glutathione metabolism in mice. *Molecular Systems Biology*, 11(10), 2015.

46. L. Zhu, S. S. Baker, C. Gill, W. Liu, R. Alkhouri, R. D. Baker, and S. R. Gill. Characterization of gut microbiomes in nonalcoholic steatohepatitis (NASH) patients: A connection between endogenous alcohol and NASH. *Hepatology*, 57(2), 2013.

47. L. Mancabelli, C. Milani, G. A. Lugli, F. Turroni, M. Mangifesta, A. Viappiani, A. Ticinesi, A. Nouvenne, T. Meschi, D. Van Sinderen, and M. Ventura. Unveiling the gut microbiota composition and functionality associated with constipation through metagenomic analyses. *Scientific Reports*, 7 (9879), 2017.

48. D. J. Morrison and T. Preston. Formation of short chain fatty acids by the gut microbiota and their impact on human metabolism. *Gut Microbes*, 7(3), 2016.

49. C. Serena, V. Ceperuelo-Mallafré, N. Keiran, M. I. Queipo-Ortuño, R. Bernal, R. Gomez-Huelgas, M. Urpi-Sarda, M. Sabater, V. Pérez-Brocal, C. Andrés-Lacueva, A. Moya, F. J. Tinahones, J. M. Fernández-Real, J. Vendrell, and S. Fernández-Veledo. Elevated circulating levels of succinate in human obesity are linked to specific gut microbiota. *ISME Journal*, 12(7), 2018.

50. S. Magnúsdóttir, D. Ravcheev, V. De Crécy-Lagard, and I. Thiele. Systematic genome assessment of B-vitamin biosynthesis suggests cooperation among gut microbes. *Frontiers in Genetics*, 6(148), 2015.

51. V. Sharma, D. A. Rodionov, S. A. Leyn, D. Tran, S. N. Iablokov, H. Ding, D. A. Peterson, A. L. Osterman, and S. N. Peterson. B-Vitamin Sharing Promotes Stability of Gut Microbial Communities. *Frontiers in Microbiology*, 10(1485), 2019.

52. H. Qi, Y. Li, H. Yun, T. Zhang, Y. Huang, J. Zhou, H. Yan, J. Wei, Y. Liu, Z. Zhang, Y. Gao, Y. Che, X. Su, D. Zhu, Y. Zhang, J. Zhong, and R. Yang. Lactobacillus maintains healthy gut mucosa by producing L-Ornithine. *Communications Biology*, 2(1), 2019.

53. N. E. Lewis, K. K. Hixson, T. M. Conrad, J. A. Lerman, P. Charusanti, A. D. Polpitiya, J. N. Adkins, G. Schramm, S. O. Purvine, D. Lopez-Ferrer, K. K. Weitz, R. Eils, R. König, R. D. Smith, and B. Palsson. Omic data from evolved E. coli are consistent with computed optimal growth from genome-scale models. *Molecular Systems Biology*, 6(390), 2010.

54. D. Machado and M. Herrgård. Systematic Evaluation of Methods for Integration of Transcriptomic Data into Constraint-Based Models of Metabolism. *PLOS Computational Biology*, 10(4), 2014.

55. J. Schellenberger, N. E. Lewis, and B. Palsson. Elimination of thermodynamically infeasible loops in steady-state metabolic models. *Biophysical Journal*, 100(3), 2011.

56. F. Cazals and C. Karande. A note on the problem of reporting maximal cliques. *Theoretical Computer Science*, 407(1-3), 2008.

57. A. M. Feist, C. S. Henry, J. L. Reed, M. Krummenacker, A. R. Joyce, P. D. Karp, L. J. Broadbelt, V. Hatzimanikatis, and B. Palsson. A genome-scale metabolic reconstruction for Escherichia coli K-12 MG1655 that accounts for 1260 ORFs and thermodynamic information. *Molecular Systems Biology*, 3(121), 2007.

58. I. Thiele, T. D. Vo, N. D. Price, and B. Ø. Palsson. Expanded metabolic reconstruction of Helicobacter pylori (iIT341 GSM/GPR): an in silico genome-scale characterization of single- and double-deletion mutants. *Journal of Bacteriology*, 187(16), 2005.

59. M. L. Mo, B. Ø. Palsson, and M. J. Herrgård. Connecting extracellular metabolomic measurements to intracellular flux states in yeast. *BMC Systems Biology*, 3(37), 2009.

60. H. Hefzi, K. S. Ang, M. Hanscho, A. Bordbar, D. Ruckerbauer, M. Lakshmanan, C. A. Orellana, D. Baycin-Hizal, Y. Huang, D. Ley, V. S. Martinez, S. Kyriakopoulos, N. E. Jiménez, D. C. Zielinski, L. E. Quek, T. Wulff, J. Arnsdorf, S. Li, J. S. Lee, G. Paglia, N. Loira, P. N. Spahn, L. E. Pedersen, J. M. Gutierrez, Z. A. King, A. M. Lund, H. Nagarajan, A. Thomas, A. M. Abdel-Haleem, J. Zanghellini, H. F. Kildegaard, B. G. Voldborg, Z. P. Gerdtzen, M. J. Betenbaugh, B. O. Palsson, M. R. Andersen, L. K. Nielsen, N. Borth, D. Y. Lee, and N. E. Lewis. A Consensus Genome-scale Reconstruction of Chinese Hamster Ovary Cell Metabolism. *Cell Systems*, 3(5), 2016.

61. E. Brunk, S. Sahoo, D. C. Zielinski, A. Altunkaya, A. Dräger, N. Mih, F. Gatto, A. Nilsson, G. A. Preciat Gonzalez, M. K. Aurich, A. Prlic, A. Sastry, A. D. Danielsdottir, A. Heinken, A. Noronha,

P. W. Rose, S. K. Burley, R. M. Fleming, J. Nielsen, I. Thiele, and B. O. Palsson. Recon3D enables a three-dimensional view of gene variation in human metabolism. *Nature Biotechnology*, 36(3), 2018.

62. Z. A. King, A. Dräger, A. Ebrahim, N. Sonnenschein, N. E. Lewis, and B. O. Palsson. Escher: A Web Application for Building, Sharing, and Embedding Data-Rich Visualizations of Biological Pathways. *PLOS Computational Biology*, 11(8), 2015.

63. A. Heinken and I. Thiele. Systematic prediction of health-relevant human-microbial co-metabolism through a computational framework. *Gut Microbes*, 6(2), 2015.

64. A. Noronha, J. Modamio, Y. Jarosz, E. Guerard, N. Sompairac, G. Preciat, A. D. Daníelsdóttir, M. Krecke, D. Merten, H. S. Haraldsdóttir, A. Heinken, L. Heirendt, S. Magnúsdóttir, D. A. Ravcheev, S. Sahoo, P. Gawron, L. Friscioni, B. Garcia, M. Prendergast, A. Puente, M. Rodrigues, A. Roy, M. Rouquaya, L. Wiltgen, A. Žagare, E. John, M. Krueger, I. Kuperstein, A. Zinovyev, R. Schneider, R. M. Fleming, and I. Thiele. The Virtual Metabolic Human database: Integrating human and gut microbiome metabolism with nutrition and disease. *Nucleic Acids Research*, 47(D1), 2019.

65. A. Ebrahim, J. A. Lerman, B. O. Palsson, and D. R. Hyduke. COBRApy: COnstraints-Based Reconstruction and Analysis for Python. *BMC Systems Biology*, 7(74), 2013.

Supplementary information

Table S1. Models used for benchmarking with number of metabolites and reactions before (and after) preprocessing.

Organism	Model	Reactions	Metabolites
E. coli	e_coli_core	95 (137)	72 (72)
H. pylori	iIT341	554 (600)	485 (422)
S. cerevisiae	iMM904	1,577 (1,273)	1,226 (722)
E. coli	iML1515	2,712 (2,868)	1,877 (1,633)
C. griseus	iCHOv1	6,663 (6,805)	4,456 (2,749)
H. sapiens	Recon3D	10,600 (14,119)	5,835 (5,835)

Table S2. Models used in host-microbe model with number of metabolites and reactions. Identifiers of microbial models are from Magnúsdóttir et al.[43].

Organism	Model	Reactions	Metabolites
S. thermophilus	LMG_18311	954	851
F. prausnitzii	A2_165	1,225	1,047
L. lactis	subsp_lactis_WCFS1	1,295	1,118
L. plantarum	Il1403	1,753	1,348
E. coli	str_K_12_substr_MG1655	2,313	1,625
B. thetaiotaomicron	VPI_5482	2,524	1,549
H. sapiens	Recon3D	10,600	5,835

Algorithm S1: Finding a new viable cut set.

Input: Set of known MPS Π and set of known MCSs Γ.

Output: New viable cut set $\mathcal{C}$ and set of known MCSs Γ.

repeat

    $\mathcal{C} \leftarrow \text{BIP}(\Pi, \Gamma)$;

    Deactivate all $i \in \mathcal{C}$;

    $\mathbf{r} \leftarrow \text{LP}(\emptyset)$;

    if $\mathbf{r}$ *is not optimal* **then**

        Reactivate all $i \in \mathcal{C}$;

        $\Gamma \leftarrow \Gamma \cup \{\mathcal{C}\}$;

until *BIP is infeasible or* $\mathcal{C}$ *is not an MCS*;

Algorithm S2: Finding a new MP with iterative minimization.

Input: Subnetwork $\mathcal{S}$ and whether to randomize.

Output: New MP $\mathcal{P}$.

$\mathcal{P} \leftarrow \emptyset$;

while $\mathcal{S} \neq \emptyset$ **do**

    Deactivate a (random) $i \in \mathcal{S}$;

    $\mathcal{S} \leftarrow \mathcal{S} \setminus \{i\}$;

    $\mathbf{r} \leftarrow \text{LP}(\mathcal{S})$;

    if $\mathbf{r}$ *is optimal* **then**

        if *randomize* **then**

            Draw a random $j \in \mathcal{S}$;

            while $r_j = 0$ **do**

                Deactivate j;

                $\mathcal{S} \leftarrow \mathcal{S} \setminus \{j\}$;

                Draw a random $j \in \mathcal{S}$;

        else

            foreach $j \in \mathcal{S}$ **do**

                if $r_j = 0$ **then**

                    Deactivate j;

                    $\mathcal{S} \leftarrow \mathcal{S} \setminus \{j\}$;

    else

        Reactivate i;

        $\mathcal{P} \leftarrow \mathcal{P} \cup \{i\}$;

Algorithm S3: Enumeration with iterative minimization and graph.

Input: Subnetwork $\mathcal{S}$.
Output: Set of MPs Π and set of MCSs Γ.
$\Pi \leftarrow \Gamma \leftarrow \mathcal{Q} \leftarrow \mathcal{G} \leftarrow \emptyset$;
$\mathcal{S}_0 \leftarrow \mathcal{S}$;
repeat
    if $\mathcal{Q}$ **then**
        Draw $\mathcal{S}$ from $\mathcal{Q}$;
        $\mathcal{Q} \leftarrow \mathcal{Q} \setminus \{\mathcal{S}\}$;
    else
        $\mathcal{S} \leftarrow \mathcal{S}_0$;
    repeat
        $\mathcal{C}, \Gamma \leftarrow$ **Algorithm S1**;
        if $\mathcal{C}$ *is viable* **then**
            $\mathcal{P} \leftarrow$ **Algorithm S2**;
            $\Pi \leftarrow \Pi \cup \{\mathcal{P}\}$;
            if $\mathcal{S} = \mathcal{S}_0$ **then**
                $\mathcal{G} \leftarrow \mathcal{G} \cup \binom{\mathcal{P}}{2}$;
                Find new maximal cliques $\mathcal{K}$ in $\mathcal{G}$;
                $\mathcal{Q} \leftarrow \mathcal{Q} \cup \mathcal{K}$;
    until *no more viable cut sets or ($\mathcal{S} = \mathcal{S}_0$ and $\mathcal{Q}$)*;
until $\mathcal{S} = \mathcal{S}_0$ *and BIP is infeasible*;

Figure S1. Running time Δt_0 for finding the first new MP with direct (blue) and iterative minimization with (purple) and without (red) randomization for different subnetwork sizes $|\mathcal{S}|$ in six different models with network size $|\mathcal{N}|$. Stars indicate significant difference between means from one-way ANOVA where data was available for all three methods or t-test where data was available for two methods (significance level 0.05). Optimization was stopped after one hour (indicated by dashed line).

Figure S2. (**A**) Number of minimizations that timed out before the first MP was found and (**B**) number of enumerations completed by one method but not by any of the others by method and model.

Figure S3. MPs found in complete enumerations. (**A**) Number of MPs $|\Pi|$ as a function of subnetwork size $|\mathcal{S}|$ for enumerations completed within one hour by at least one method. Lines indicate means with 95% confidence intervals from bootstrapping and colors indicate models. (**B**) Number of MCSs $|\Gamma|$ as a function of $|\Pi|$ with colors indicating models and a dashed line indicating $|\Gamma| = |\Pi|$.

Figure S4. Number of MPs $|\Pi|$ found in one hour with direct minimization (red), iterative minimization (blue), iterative minimization with randomization (purple), and iterative minimization with graph (orange) for different subnetwork sizes $|\mathcal{S}|$ in six different models with network size $|\mathcal{N}|$. Stars indicate significant difference between means from one-way ANOVA where data was available for all three methods or t-test where data was available for two methods (significance level 0.05). Enumerations that were completed within one hour are not included.

Figure S5. Correlation of gene importance in *E. coli* central carbon metabolism predicted by MP sampling (red) and FBA (blue) with experimental growth and gene essentiality data from the Keio collection[32]. (**A**) Pearson correlation (*r*) of predicted gene importance with experimental growth (OD of knockout strain at 24 h). Fitted lines are shown with 95% confidence intervals from bootstrapping. (**B**) Spearman correlation (ρ) of predicted gene importance with experimental gene essentiality score (higher score means more evidence for essentiality and vice versa). Means are indicated with 95% confidence intervals from bootstrapping.