An Efficient Implementation of Flux Variability Analysis for Metabolic Networks

Bruno G. Galuzzi[1,2(✉)] and Chiara Damiani[1,2]

[1] Department of Biotechnology and Biosciences, University of Milano-Bicocca, Milan, Italy
{bruno.galuzzi,chiara.damiani}@unimib.it
[2] SYSBIO Centre of Systems Biology/ ISBE.IT, Milan, Italy

Abstract. Flux Variability Analysis (FVA) is an important method to analyze the range of fluxes of a metabolic network. FVA consists in performing a large number of independent optimization problems, to obtain the maximum and minimum flux through each reaction in the network. Although several strategies to make the computation more efficient have been proposed, the computation time of an FVA can still be limiting. We present a two-step procedure to accelerate the FVA computational time that exploits the large presence within metabolic networks of sets of reactions that necessarily have an identical optimal flux value or only differ by a multiplication constant. The first step identifies such sets of reactions. The second step computes the maximum and minimum flux value for just one element of each of set, reducing the total number of optimization problems compared to the classical FVA. We show that, when applied to any metabolic network model included in the BiGG database, our FVA algorithm reduces the total number of optimization problems of about 35%, and the computation time of FVA of about 30%.

Keywords: Metabolic networks · Flux Balance Analysis · Flux variability Analysis · Constrained-based modeling

1 Introduction

The study of cell metabolism is of paramount importance in various fields, including health, wellness, and bio-transformations [4]. Indeed, metabolism is related with most cellular processes and may act as an integrative readout of the pato-physiological state of a cell [18]. The first requirement to understand metabolism is the knowledge of metabolic fluxes, i.e. the velocities and the directions of all the biochemical reactions involved in all metabolic processes, such as glycolysis and oxidative phosphorylation.

Currently, direct determination of metabolic fluxes at the genome-wide level is not feasible. However, they can be predicted numerically, via integration of multiple -omics data (e.g., transcriptomics, proteomics, and metabolomics) into

constraint-based stoichiometric metabolic models [3,4,13]. The starting point of constrained-based modeling is the information embedded in the metabolic network, which represents the set of all the possible biochemical reactions that can occur in a cell in a specific organism or tissue. Such information can be represented with a stoichiometric matrix S of dimension $m \times n$, where m is the number of metabolites and n is the number of reactions. In a constrained-based metabolic model, a steady-state condition is imposed, that is, the total production of any metabolite must equal to the total amount of its consumption. Hence, a possible metabolic flux configuration is represented by a vector $\vec{v}$, for which $S\vec{v} = 0$, i.e. the null space of the stoichiometric matrix. Additional constraints, such as thermodynamics or capacity constraints, can also be incorporated.

Despite the large number of constraints, a metabolic network typically includes more reactions than metabolites, resulting in a large space of possible feasible flux distributions. In this case, different strategies can be used to predict the target metabolic flux distribution. A possibility consists in using efficient flux sampling strategies [7,12] to explore the entire region of feasible flux solutions and obtaining information on the range of feasible flux solutions and on their probabilities. Another possibility is Flux Balance Analysis (FBA), which assumes that the cell behaviour is optimal with respect to an objective function, such as the biomass production rate. The objective function defines a reaction that must be maximized or minimized under the set constraints. FBA identifies a single solution among the set of possibly many alternative optimal solutions.

Alternatively, it is possible to study the range of each metabolic flux across the set of alternative optima by means of Flux Variability Analysis (FVA) [16]. In a nutshell, FVA consists in finding the minimum and maximum flux through each reaction in the network, given some constraints on the state of the network, e.g., imposing a minimum percentage of the maximal biomass production rate. Typical applications of FVA in systems biology include investigating network flexibility and network redundancy [24]. Recently, FVA has gained importance as a preliminary step for omics data integration. For example, in [5,8], to sensibly limit the flux of an internal reaction based on gene expression data, we first needed to compute the maximum and minimum flux through such reaction, based on nutrient availability constraints. In fact, limiting the flux boundaries defined *a priori* by the modeler might produce no effect on the flux if they are larger than the actual maximal flux determined by the environment.

Given a generic metabolic network with n reactions, FVA requires $2n$ optimization problems to be solved, maximizing and minimizing the flux of any reaction in the model. Given the high number of possible reactions and constraints characterizing a metabolic network, efficient computation of FVA is fundamental. The aim of this work consists of exploring a new way to partially reduce the computational time classically required to solve FVA.

1.1 State of the Art

The computation of FVA clearly can be distributed and is therefore ideally suited for high performance computing. For example, the current FVA implementations

in Cobrapy [6], COBRA toolbox [22], and COBRAjl [11] split the optimization problem into multiple processes. Other implementation have been designed to reduce the computational time requested to find an initial feasible solution. For example, fastFVA [23], one of the most used efficient implementations of FVA, coded in C++, iterates through all the reactions and solves the two optimization problems using a specialized LP solver, such as GLPK or CPLEX, but without spending time effort in finding a feasible solution or pre-processing the linear system, since the feasible region is the same for all the optimization problems. Finally, implementations of FVA exist that include thermodynamic constraints [17] to remove unbounded fluxes through reactions contained in internal cycles.

1.2   Our Contribution

To the best of our knowledge, no current implementation of FVA exploits the fact that in a generic metabolic network, there are a lot of constraints of the form $av_i + bv_j = 0$. Hence, if $v_i^{\max}$ and $v_i^{\min}$ have already been computed, one can omit the two optimization problems for $v_j^{\max}$ and $v_j^{\min}$. This simple but effective consideration is expected to considerably reduce the total number of optimization problems for FVA of any metabolic network. To clarify this observation with an example, in Fig. 1, we show a small sub-network of the human metabolism related to the glycolysis and Pentose phosphate pathways, represented using the web-app Escher [14]. Each blue arrow represents a reaction and each node represents a metabolite involved in a reaction. Under the steady-state assumption, the total production of any metabolite must be equal to the total amount of its consumption. It can be noticed that many metabolites, such as 3-phospho-D-glycerate (3pg_c), D-glycerate 2-phosphate (2pg_c), and phosphoenolpyruvate (pep_c) are involved in two different reactions only. Hence, the steady-state assumption imposes that the flux through phosphoglycerate kinase (PGK) and phosphoglycerate mutase (PGM) must be equal, and that the flux of
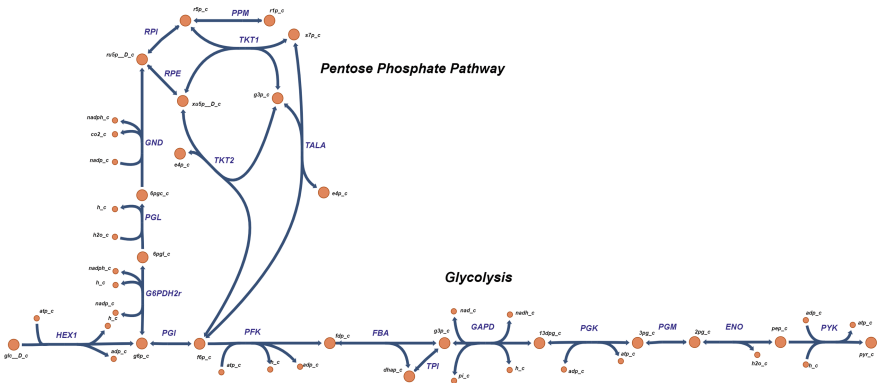

Fig. 1. Example of human metabolic sub-network related to glycolysis and Pentose phosphate pathways.

the Phosphoglycerate kinase (PGK) and Enolase (ENO) must be equal. Hence, once the maximum and minimum flux through the PGK reaction has been found, one can avoid to solve the optimization problems for the PGK and ENO reactions. In this work, we investigated the presence of these type of constraints involving only two fluxes in all the metabolic models publicly available in the BiGG model database [19]. Then, we used such information to implement a Python-based Cobrapy [6] extension that improves the computation of FVA including an efficient pre-FVA step to find all the possible reactions for which the optimization problems can be omitted. Such pre-FVA step can be used not only for FVA but also for all computational tools requiring the computation of several optimization problems for the same metabolic networks, such as finding blocked reactions, searching for essential reactions of a target objective function, and reaction deletion analysis.

2   Material and Methods

2.1   COBRA Model

Assuming that cell behavior is optimal with respect to an objective, optimization methods, such as Flux Balance Analysis (FBA) [21], can be used to calculate an optimal flux distribution with respect to a specific objective. In a nutshell, a metabolic network is associated with the following linear programming (LP) problem:

$$\max \sum_{i=1}^{r} w_i v_i \qquad (1)$$
$$S \cdot \vec{v} = \vec{0}$$
$$\vec{v_L} \leq \vec{v} \leq \vec{v_U}$$

where w_i is the objective coefficient of flux i, and $\vec{v_L}$ and $\vec{v_U}$ represent the possible bounds used to mimic as closely as possible the biological process in the analysis.

2.2   Flux Variability Analysis

Flux Variability Analysis (FVA) [9,16] is a constraint-based modeling technique aimed at determining the maximal (and minimal) possible flux through any reaction of the model, to evaluate the cell's range of metabolic capabilities. FVA solves the following two linear programming optimization problems (one for minimization and one for maximization) for each flux v_j of interest, with $j = 1, \ldots, n$:

$$\max / \min \ v_j \qquad (2)$$
$$S \cdot \vec{v} = \vec{0}$$
$$\vec{v_L} \leq \vec{v} \leq \vec{v_U}$$
$$\vec{v_i} \geq \gamma Z_0$$

where Z_0 is an optimal solution for Eq. 1, and γ is a parameter, which controls whether the analysis is done w.r.t. suboptimal network states ($0 \leq \gamma < 1$) or to the optimal state ($\gamma = 1$). FVA can be used also to search for blocked reactions in a network. A blocked reaction r_i in a metabolic network is a reaction that carries no flux in any feasible solution (i.e. $v_i = 0$). Practically, the blocked reactions are the subset of reactions for which $\max v_j = \min v_j = 0$.

2.3   Pre-step to Accelerate FVA

Our improvement of FVA computation is based on a pre-FVA operation to select the subset of reactions for which it is really necessary to solve the maximization and the minimization problems. We showed this step in Fig. 2. First, we considered the set of all the possible reactions $\{v_1, \ldots, v_n\}$, and we split the reactions in connected-subsets. Each connected-subset is formed by reactions whose flux value can be derived by multiplying any other flux value of the subset by a known constant (e.g., $2v_1 - 3v_2 = 0$), which may also take value 1 (e.g., $v_1 - v_2 = 0$). For the computation of the connected-subsets, we built a graph whose nodes represent the reactions $r_1, \ldots, r_n$ of the network. We connected two nodes r_i and r_j if and only if a constrain of the form $av_i + bv_j = 0$ exists. The connected-subsets corresponds with the connected components of such graph. Then, we created a set I formed by one arbitrary reaction for each of these subsets, and we performed FVA for the reactions in I only. Afterwards, we derived the rest of the FVA values using the direct relation with the FVA values computed before. This implies to solve a linear system in which the variables are the fluxes of the reactions not belonging to I.
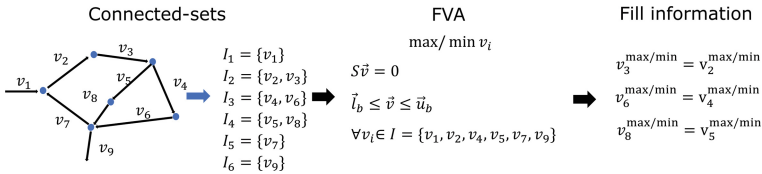

Fig. 2. Schematization of the main steps involved in the computation of efficient FVA.

2.4   Datasets

To demonstrate the applicability of our strategy to real-world metabolic networks, and the improvement as compared to classic FVA, we considered all the metabolic networks in the BiGG database [19]. BiGG Models integrates many published genome-scale metabolic networks into a single database with a set of standardized identifiers. BiGG models include information on the metabolic genes associated with reactions in the form of Gene Protein Reaction rules. Metabolites are linked to many external databases (e.g. KEGG, PubChem).

The total number of metabolic networks is 108 collected from 85 different organisms. The dimension of the models span from 95 reactions and 72 metabolites (e_coli_core [20]) to 10,600 reactions and 5,835 metabolites (Recon3D [2]).

2.5   Software Availability and Computational Architecture

To analyze the COBRA models and perform classical FVA, we used the functions provided by the Cobrapy toolkit [6]. To compute the connection-sets, we wrote a specific Python code based on the function *connected_components* provided by the Scipy library [26]. The code can be promptly integrated into the Cobrapy toolkit. All computations were performed on an Intel(R)@3GHz 32GB, using Gurobi as solver and one single CPU.

The source code and documentation are available at https://github.com/CompBtBs/efficientFVA.

3   Experimental Results

3.1   Metabolic Networks Present Many Connected Sets

We investigated the number of connected-sets of each metabolic network in the BiGG database. In particular, we computed the number of reactions, the number of connected-sets, the number of connected-sets formed by 2 reaction at least, and the dimension of the largest connected-set. In Table 1, we reported the results for 10 of the 108 networks. We also reported in Fig. 3a a histogram of the distribution of the ratio between the number of connected-sets and the number of reactions for all the 108 models.

At first instance, all the networks have less connected-sets than the number of reactions. The ratio between the total number of connected-subsets and the reactions ranges between 0.48 (iJB785 [1]) and 0.79 (iLB1027_lipid [15]), with mean 0.63 and standard deviation 0.04. This indicates that, in any network, there are a lot of constraints of the form $av_i + bv_j = 0$. Moreover, some models include very large connected-sets. For example, in the iCHOv1 model [10], the largest connected-set is made by 56 reactions. The presence of these connected-sets is due to the intrinsic nature of metabolic networks. Indeed, many metabolites are involved in just two reactions: one reaction consumes and the other produces such metabolite. Moreover, linear chains of reactions exist in which a series of reactions are connected linearly to produce/consume specific metabolites involved in just two of these reactions. From this analysis, we found that the total number of optimizations required for the FVA can be reduced by at least 35% in most of the networks, when considering a single reaction per connected set.

Table 1. Number of reactions, connected-sets (CS), connected-sets formed by 2 reaction at least (2-CS), and dimension of the largest connected-set (max-CS), for ten networks from the BiGG database.

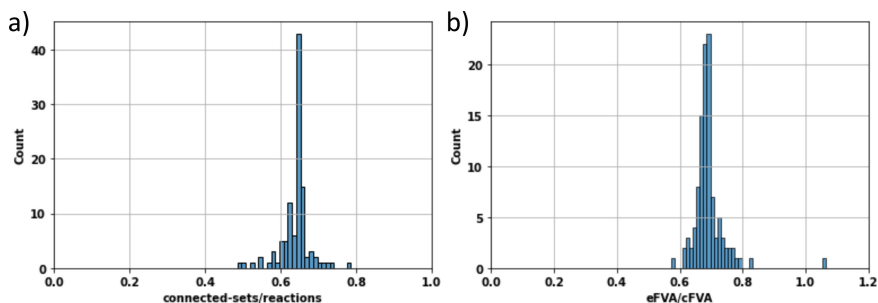| Model	Reactions	CS	CS /Reactions	2-CS	max-CS
e_coli_core	95	62	0.65	26	4
iAB_RBC_283	469	291	0.62	98	16
iIS312_Amastigote	519	321	0.62	117	14
iAT_PLT_636	1008	645	0.64	176	21
iRC1080	2191	1620	0.74	298	26
iYL1228	2262	1353	0.60	512	15
iMM1415	3726	2346	0.63	668	38
RECON1	3741	2329	0.62	648	38
iCHOv1	6663	4793	0.72	927	56
Recon3D	10600	7414	0.70	1697	28

Fig. 3. a) Histogram of the distribution of the ratio between the number of connected-sets and the number of reactions for all the 108 models. b) Histogram of the distribution of the ratio between the mean computational time of our implementation of FVA (eFVA) and of classical FVA (cFVA), for all the 108 models.

3.2 Connected-Sets Reduce the Computational Time for FVA

We investigated the possible computational savings that can be achieved by computing the connected-sets of a metabolic network. In Table 2, we reported: the mean ± the standard deviation over five different runs of the computational time of classical FVA (cFVA), of our overall implementation (eFVA), and of the step for the identification of the connected-sets (pre-FVA) alone. Note that the reported computational time of eFVA includes the pre-FVA computational time. The table also reports the ratio between the mean computational time of

our implementation of FVA(eFVA) and of the classical FVA(cFVA). We also reported in Fig. 3b an histogram of the ratio between eFVA and cFVA for all the 108 models.

Table 2. Computational time for cFVA, eFVA, pre-FVA step, and the ratio between eFVA and cFVA on ten metabolic networks from the BiGG database.

Model	cFVA(s)	eFVA(s)	pre-FVA step(s)	eFVA/FVA
e_coli_core	0.06 ± 0.0	0.07 ± 0.02	0.02 ± 0.0	1.06
iAB_RBC_283	0.63 ± 0.01	0.46 ± 0.0	0.06 ± 0.0	0.72
iIS312_Amastigote	0.83 ± 0.01	0.6 ± 0.01	0.07 ± 0.01	0.72
iAT_PLT_636	3.82 ± 0.02	2.58 ± 0.03	0.14 ± 0	0.67
iRC1080	15.72 ± 0.08	12.37 ± 0.1	0.31 ± 0.01	0.79
iYL1228	16.1 ± 0.04	10.99 ± 0.12	0.72 ± 0.05	0.68
iMM1415	53 ± 0.32	37.36 ± 0.24	2.35 ± 0.01	0.70
RECON1	41.16 ± 0.06	28.64 ± 0.05	2.27 ± 0.01	0.70
iCHOv1	150.08 ± 0.18	116.47 ± 0.21	6.1 ± 0.16	0.78
Recon3D	246.03 ± 0.24	204.11 ± 0.38	30.62 ± 0.06	0.83

In all the cases, we observed a standard deviation for cFVA, eFVA, and the pre-FVA step of at least one order of magnitude less then the corresponding mean values. The computational time required for eFVA results less than cFVA in all the cases, except for the smallest network (e_coli_core), for which the time required for cFVA and eFVA results similar. This is due to the additional time required to the computation of the connected-subsets that, for this case, represents about 30% of the entire eFVA time. For all the other cases, the time required for the pre-FVA step results one or more order of magnitude less than the entire eFVA time. More importantly, the ratio between the eFVA and cFVA ranges between 0.57 (iJB785) and 0.83 (Recon3D) with mean 0.69 and standard deviation 0.05.

We expected the number of metabolites, involved in just two reactions, to increase with the number of reactions, and hence the gain of using eFVA to be more evident for larger metabolic networks. Indeed, the Pearson correlation between the ratio of the number of connected-sets over the number of reactions and the ratio of the mean computational time of eFVA over cFVA, for all the 108 models is significant, namely 0.49 ($p_{value} < 0.001$), 0.67 ($p_{value} < 0.001$) if the outlier corresponding with smallest network (e_coli_core) is removed. The larger gain for larger networks can also be noticed in Fig. 4, where we reported the time required for cFVA, eFVA, and pre-FVA step, as a function of the number of reactions.
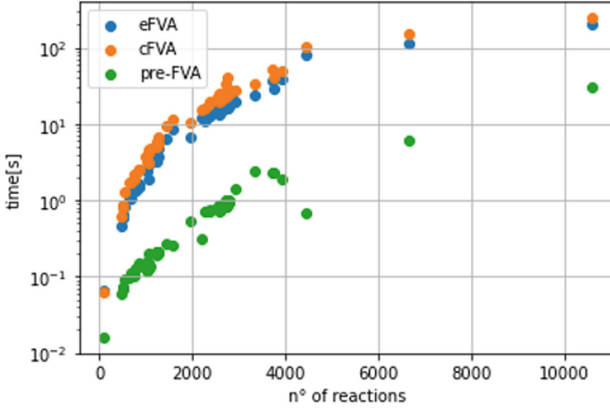
Fig. 4. Scatter-plot showing the time required for cFVA, eFVA, and pre-FVA step, as a function of the number of reactions.

3.3 Connected-Sets Improve the Search of Blocked Reactions in Large Metabolic Networks

The information of the connected-sets can be used not only for FVA, but also to accelerate the search for blocked reactions. To investigate this fact, we computed the subset of blocked reactions for the ten metabolic networks analyzed in Table 1. In Table 3, we reported the computational time required to search blocked reactions using the implementation provided by the Cobrapy library (cBR) and our version based on the pre-computation of connected-sets (eBR). Again, we reported the mean and standard deviation of five different runs. Surprisingly, all the networks, except iAT_PLT_636 [25] and Recon3D, show a non negligible number of blocked reactions varying between 4.31% to 59.92%. The

Table 3. Number of blocked reactions, and computational time for cBR, eBR, and pre-FVA step and the ratio between eFVAand cFVAon ten metabolic networks from the BiGG database.

Model	% blocked reactions	cBR(s)	eBR(s)	eBR/cBR
e_coli_core	8.42	0.02 ± 0.00	0.02 ± 0.01	1.44
iAB_RBC_283	3.41	0.19 ± 0.01	0.19 ± 0.06	1.02
iIS312_Amastigote	59.92	0.17 ± 0.01	0.18 ± 0.08	1.10
iAT_PLT_636	0.00	1.99 ± 0.08	1.46 ± 0.14	0.73
iRC1080	27.75	8.04 ± 0.24	6.62 ± 0.31	0.82
iYL1228	45.93	7.06 ± 0.02	5.51 ± 0.79	0.78
iMM1415	34.73	43.02 ± 2.18	34.9 ± 2.49	0.81
RECON1	34.06	36.65 ± 0.69	27.42 ± 2.48	0.75
iCHOv1	35.76	135.4 ± 17.66	121.36 ± 6.27	0.90
Recon3D	0.00	333.55 ± 37.57	275.48 ± 31.64	0.83

computational time required for eBR results less than cBR in seven out of ten networks. For moderate sizes of the metabolic network ($r < 1000$), the use of connected-sets does not improve or even worsen (e_coli_core) the time required to find the blocked reactions. Note that, for the computation of blocked reactions, the Cobrapy implementation has a pre-processing step in which a general optimization is solved setting randomly a target function. Then, all the reactions for which the fluxes of the optimal solution differ from 0 era excluded from the $2n$ optimization problems. This pre-processing step already speeds up the computation of blocked reactions, and so the time saving by our implementation results less important for small metabolic networks. On the contrary, for metabolic networks having more than 1000 reactions, the time required for the computation of the connected sets results one or more order of magnitude less than the entire eBR time and the ratio between the eBR and cBR results between 0.73 (iAR_PLT_636) and 0.83 (Recon3D).

4   Conclusions

The numerical computation of plausible metabolic fluxes using constrained-based modeling is acquiring increasingly relevance to understand the mechanisms related to the physio-pathological state of a cell or an organism. To this aim, efficient tools for the analysis of large genome-wide metabolic networks are mandatory. In this work, we have considered the computation of Flux Variability Analysis, and proposed a simple but effective method to accelerate such a computation by means of the connected-sets formed by reactions whose flux value can be directly derived from the flux value of any reaction in the same set. Of course, our modification of classical FVA is useful only when the number of connected-sets is less than the total number of reactions. However we verified that this fact holds for all the networks in the BiGG database. In our experiments, we were able to reduce the total number of optimization problems of about 35%, and the computation time of FVA of about 30%. Obviously, the quantity of saved computational time depends intrinsically on the quantity of constraints of the form $av_i + bv_j = 0$ in the network. However, we have shown that this quantity is non negligible in all current metabolic network reconstructions. Therefore, even considering the time required for the computation of the connected-sets in the overall computation time, our approach significantly improves the efficiency in practise.

In this work, to compute the flux variability of the reactions for which it necessarily must be computed after the identification of connected-sets, we relied on the FVA implementation in COBRApy, and used a single CPU. However, we remark that our implementation does not need a specific implementation of FVA for this step. Therefore, to further improve the computation process, one could use our approach in combination with either parallelization or more efficient versions of FVA, such as the one proposed in Thiele and Gudmundsson [23].

As a further work, we propose to exploit other possible specific constraints coming e.g. from the integration of -omics data, to reduce even more the total number of optimization problems necessary for FVA.

References

1. Broddrick, J.T., et al.: Unique attributes of cyanobacterial metabolism revealed by improved genome-scale metabolic modeling and essential gene analysis. Proc. Natl. Acad. Sci. **113**(51), E8344–E8353 (2016)

2. Brunk, E., et al.: Recon3d enables a three-dimensional view of gene variation in human metabolism. Nat. Biotechnol. **36**(3), 272–281 (2018)

3. Cavill, R., Jennen, D., Kleinjans, J., Briedé, J.J.: Transcriptomic and metabolomic data integration. Briefings Bioinf. **17**(5), 891–901 (2016)

4. Damiani, C., Gaglio, D., Sacco, E., Alberghina, L., Vanoni, M.: Systems metabolomics: from metabolomic snapshots to design principles. Curr. Opin. Biotechnol. **63**, 190–199 (2020)

5. Di Filippo, M., et al.: Integrate: model-based multi-omics data integration to characterize multi-level metabolic regulation. PLoS Comput. Biol. **18**(2), e1009337 (2022)

6. Ebrahim, A., Lerman, J.A., Palsson, B.O., Hyduke, D.R.: COBRapy: constraints-based reconstruction and analysis for python. BMC Syst. Biol. **7**(1), 1–6 (2013)

7. Fallahi, S., Skaug, H.J., Alendal, G.: A comparison of monte Carlo sampling methods for metabolic network models. PLoS ONE **15**(7), e0235393 (2020)

8. Galuzzi, B.G., Vanoni, M., Damiani, C.: Combining denoising of RNA-seq data and flux balance analysis for cluster analysis of single cells. BMC Bioinform. **23**(6), 1–21 (2022)

9. Gudmundsson, S., Thiele, I.: Computationally efficient flux variability analysis. BMC Bioinform. **11**(1), 1–3 (2010)

10. Hefzi, H., et al.: A consensus genome-scale reconstruction of Chinese hamster ovary cell metabolism. Cell Syst. **3**(5), 434–443 (2016)

11. Heirendt, L., Thiele, I., Fleming, R.M.: DistributedFBA. jl: high-level, high-performance flux balance analysis in Julia. Bioinformatics **33**(9), 1421–1423 (2017)

12. Herrmann, H.A., Dyson, B.C., Vass, L., Johnson, G.N., Schwartz, J.M.: Flux sampling is a powerful tool to study metabolism under changing environmental conditions. NPJ Syst. Biol. Appl. **5**(1), 1–8 (2019)

13. Hyduke, D.R., Lewis, N.E., Palsson, B.Ø.: Analysis of omics data with genome-scale models of metabolism. Mol. BioSyst. **9**(2), 167–174 (2013)

14. King, Z.A., Dräger, A., Ebrahim, A., Sonnenschein, N., Lewis, N.E., Palsson, B.O.: Escher: a web application for building, sharing, and embedding data-rich visualizations of biological pathways. PLoS Comput. Biol. **11**(8), e1004321 (2015)

15. Levering, J., et al.: Genome-scale model reveals metabolic basis of biomass partitioning in a model diatom. PLoS ONE **11**(5), e0155038 (2016)

16. Mahadevan, R., Schilling, C.H.: The effects of alternate optimal solutions in constraint-based genome-scale metabolic models. Metab. Eng. **5**(4), 264–276 (2003)

17. Müller, A.C., Bockmayr, A.: Fast thermodynamically constrained flux variability analysis. Bioinformatics **29**(7), 903–909 (2013)

18. Nielsen, J.: Systems biology of metabolism: a driver for developing personalized and precision medicine. Cell Metab. **25**(3), 572–579 (2017)

19. Norsigian, C.J., et al.: BiGG models 2020: multi-strain genome-scale models and expansion across the phylogenetic tree. Nucleic Acids Res. **48**(D1), D402–D406 (2020)

20. Orth, J.D., Fleming, R.M., Palsson, B.Ø.: Reconstruction and use of microbial metabolic networks: the core Escherichia coli metabolic model as an educational guide. EcoSal Plus 4(1) (2010)

21. Orth, J.D., Thiele, I., Palsson, B.Ø.: What is flux balance analysis? Nat. Biotechnol. **28**(3), 245–248 (2010)
22. Schellenberger, J., et al.: Quantitative prediction of cellular metabolism with constraint-based models: the cobra toolbox v2. 0. Nat. Protoc. **6**(9), 1290–1307 (2011)
23. Thiele, I., Gudmundsson, S.: Computationally efficient flux variability analysis. BMC Bioninform. **11**(489), 1–3 (2010)
24. Thiele, I., Fleming, R.M., Bordbar, A., Schellenberger, J., Palsson, B.Ø.: Functional characterization of alternate optimal solutions of Escherichia coli's transcriptional and translational machinery. Biophys. J . **98**(10), 2072–2081 (2010)
25. Thomas, A., Rahmanian, S., Bordbar, A., Palsson, B.Ø., Jamshidi, N.: Network reconstruction of platelet metabolism identifies metabolic signature for aspirin resistance. Sci. Rep. **4**(1), 1–10 (2014)
26. Virtanen, P., et al.: SciPy 1.0: fundamental algorithms for scientific computing in Python. Nat. Methods **17**(3), 261–272 (2020)