New genome scale network modeling and mining workflow for detecting metabolic changes induced by exposure to chemicals

Louison Fresnais, Olivier Perin, Anne Riu, Romain Grall, Alban Ott, Bernard Fromenty, Jean-Clément Gallardo, Maximilian Stingl, Clément Frainay, Fabien Jourdan, et al.

New genome scale network modeling and mining workflow for detecting metabolic changes induced by exposure to chemicals

Louison Fresnais[1,2*], Olivier Perin[2], Anne Riu[2], Romain Grall[2], Alban Ott[2], Bernard Fromenty[3], Jean-Clément Gallardo[1], Maximilian Stingl[1], Clément Frainay[1], Fabien Jourdan[1,4*] and Nathalie Poupin[1*]

1 UMR1331 Toxalim (Research Centre in Food Toxicology), Université de Toulouse, INRAE, ENVT, INP-Purpan, UPS, Toulouse, France.

2 L'Oréal Research and Innovation, Aulnay-sous-Bois, France.

3 INSERM, Univ Rennes, INRAE, Institut NUMECAN (Nutrition Metabolisms and Cancer) UMR_A 1317, UMR_S 1241, F-35000 Rennes, France

4 MetaboHUB-MetaToul, National Infrastructure of Metabolomics and Fluxomics, Toulouse, France

∗To whom correspondence should be addressed.

Abstract

The growing abundance of *in vitro* omics data, coupled with the necessity to reduce animal testing in the safety assessment of chemical compounds and even eliminate it in the evaluation of cosmetics, highlights the need for abundant computational methodologies. Data from omics technologies allow the exploration of a wide range of biological processes, therefore providing a better understanding of mechanisms of action (MoA) related to chemical exposure in biological systems. However, the analysis of these large datasets remains difficult due to the complexity of modulations spanning multiple biological processes. To address this, we propose a new computational workflow that combines knowledge on endogenous metabolism from a genome scale metabolic network (GSMN) and *in vitro* transcriptomics data with the aim of better identifying the metabolic MoA (mMoA) of chemicals. Our workflow proceeds in three main steps. The first step consists of building cell condition-specific models representing the metabolic impact of each exposure condition while taking into account the diversity of possible optimal solutions with a partial enumeration algorithm. In a second step, based on these enumerations, two conditions can be compared by extracting differentially activated reactions (DARs) between the two sets of enumerated possible condition-specific models. Finally, in the third step, DARs are grouped into clusters of functionally interconnected metabolic reactions using the distance-based clustering and subnetwork extraction method. The first part of the workflow was exemplified on eight molecules selected for their known human hepatotoxic outcomes associated with specific MoAs well described in the literature and for which we retrieved primary human hepatocytes (PHH) transcriptomic data in Open TG-GATEs. Then, we applied this new workflow to model and visualize associated mMoA for two of these eight molecules (amiodarone and valproic acid). Despite large disparities in transcriptomic effects for these two chemicals, *i.e.*, two differentially expressed genes (DEGs) for amiodarone *vs* 5709 DEGs for valproic acid, our results well fitted evidence from the literature regarding known MoA. Beyond these confirmations, the workflow highlighted potential other unexplored mMoA.

Author summary

There is an urgent need for development and validation of new approach methodologies (NAMs) to avoid animal testing in safety evaluation. Among these NAMs, the exploration of big data and the use of transcriptomics data generated from *in vitro* systems are key. Omics data reflect how cellular biological processes are globally impacted by chemical exposure, but deciphering underlying modulations remains a challenge. In particular, interactions between biological processes are not taken into account, which does not allow for the analysis of MoAs spanning several processes. To this end, we propose an original workflow able to construct condition-specific metabolic

networks from gene expression data to extract functional information about how endogenous metabolism is disrupted and visualize this mechanistic information thanks to a graph-based network analysis procedure. We highlight these new computational workflow capabilities by predicting and analyzing the metabolic impact of two known hepatotoxic compounds: amiodarone and valproic acid.

1 Introduction

Toxicology is entering a new era with the urgent need to follow a 3R (Reduce, Replace and Refine) policy when assessing risks of chemical molecules. The European Cosmetics regulation (EC) No 1223/2009 banning animal testing for cosmetic ingredients is a striking example of the need to develop non-animal approaches, particularly for systemic toxicity. In that context, new approach methodologies (NAMs) from the combination of *in silico* and *in vitro* methods are required to be fit for purpose and protective of human health. These NAMs are now being developed to support the so-called next generation risk assessment (NGRA) [1]. NAMs are already evolving at a fast pace thanks to the ever-increasing amount of omics data generated from *in vitro* experiments, creating unprecedented capacity to study biological systems. Omics screening allows, for instance, a more holistic classification of compounds based on their global effects [2,3] and it can be integrated to improve quantitative structure-activity relationships strategies [4]. The exploration of most biological processes is covered by transcriptomics or proteomics data. Among these processes, endogenous metabolism is becoming a source of concern since chemicals have the capacity to affect metabolism at a cellular and tissue level, potentially leading to adverse effects for humans such as diabetes, obesity, or even organ dysfunction [5–9]. Moreover, generating and analyzing several large-omics datasets can be expensive and methodologically challenging. Such datasets are usually of large dimensions (thousands of genes, proteins, and metabolites) that will contain information for broader insights on how an organism or a cell is globally impacted by a chemical. One of the current key challenges in the field is to extract knowledge from these rich but complex datasets.

The volumes of data generated by omics approaches and their complexity require advanced statistical and computational solutions to pinpoint patterns of interest among thousands of variables. Dimension reduction (PCA, MDS, t-SNE, etc.) methods are widely used to describe and visualize these large multidimensional datasets, but they are not designed for functional interpretation of omics data. From that perspective, enrichment-based methods, which allow the identification of biological modules (gene sets, metabolic pathways and cellular functions) in which identified modulated variables are over-represented, have the advantage of providing an overview of the studied biological modulations [10]. Nevertheless, these methods rely on arbitrary definitions of the pathways and functional sets, which might differ depending on the selected database, and tend to hide functional processes spanning several pathways [11,12]. Many studies [13–16] aim to take advantage of published gene expression data available in databases such as DrugMatrix [17], Connectivity Map [18,19], ToxicoDB [20] and Open TG-GATEs [21] (https://toxico.nibiohn.go.jp) to improve chemical toxicity assessment. For instance, Heusinkveld *et al.* [22] implemented an approach based on the comparison of Open TG-GATEs top 50 DEG signatures ranked according to their t-statistic. They aimed at providing both a score to compare compounds and a mechanistic understanding based on enrichment-based methods. On the other hand, Ting Li *et al.* [2] trained a deep neural network model for drug-induced liver injury prediction based on the LINCS L1000 dataset [19]. In both cases, metabolism is not the primary focus of the study, and the functional interpretation relies on enrichment-based methods, with the limitations explained above.

Therefore, to improve our understanding of how chemicals impact endogenous metabolism and can trigger potential adverse effects, it is necessary to develop new computational methods that would allow provision of functional

metabolic information from omics data including the diversity of possible mMoAs leading to adverse outcomes. Genome-scale metabolic networks (GSMNs) are biological networks representing all the possible biochemical reactions occurring in an organism. They are therefore well suited to consider the diversity of endogenous metabolic disruptions potentially associated with adverse outcomes at a cellular or at an organ level. These networks are reconstructed from an annotated genome, curated, and refined thanks to biochemical knowledge retrieved from the literature [23,24]. Several GSMNs representing human metabolism such as Recon2 [25], Recon2.2 [26], Recon3D [27] and Human-GEM [28] have been published. They are composed of thousands of reactions and metabolites interconnected according to the stoichiometric matrix of the network. For instance, Recon2.2, which is one of the most used human GSMNs, is composed of 7785 reactions, 5324 metabolites, and 1675 associated genes distributed over 10 cellular compartments. The stoichiometric matrix is the mathematical representation of the biochemical interdependences between reactions, containing the proportions of substrates and products involved in each reaction. The stoichiometric matrix is the ground for constraint-based modeling approaches, which are largely used to model metabolic networks at the scale of cells, tissues, or organisms. Genes and reactions are linked by gene-protein-reaction (GPR) associations, which are formulated as Boolean rules, including all genes coding for either the isoenzymes (OR association) or the enzymatic complex subunits (AND association) catalyzing a reaction. Since these networks include all the reactions that can potentially be expressed in any tissue or cell type, and any condition for a given organism, they represent the global metabolism of this organism. It is then crucial to tailor these models to specifically represent the metabolism of one tissue, one cell, or one metabolic condition in order to accurately decipher the mMoA of chemical compounds.

Preliminary to any interpretation, it is essential to perform a condition-specific metabolic network reconstruction to avoid interpreting metabolic functions that would involve reactions that are not active in the biological condition or cell type under study. Condition-specific modeling approaches aim at exploiting experimental data to assemble a GSMN that more closely represents the condition under study. Many algorithms such as iMAT [29,30], MOOMIN [31], FASTCORE [32], or INIT [33] have been developed to build condition-specific metabolic networks, most of which rely on transcriptomic data. iMAT, which is one of the most used algorithms, requires conversion of continuous gene expression values to categorical information such as “highly expressed gene” or “lowly expressed gene.” The iMAT algorithm aims to find the best consensus between the reaction activity inferred from categorized gene expression data and the activity inferred from the GSMN structure, which defines the biochemical interdependencies between reactions through consumed and produced metabolites. The output is a subnetwork (subset of all reactions of the generic metabolic network) that more faithfully represents the metabolic state of the cell in the condition of the transcriptomic experiment.

However, due to the high complexity of GSMNs and the relatively low amount of biological data, the mixed integer linear problem solved by these algorithms to find a condition-specific metabolic network is under-constrained. Practically, it means that many equally optimal condition-specific subnetworks exist for a single biological condition and that arbitrarily taking one subnetwork would limit the reproducibility and bias further analyses made for that condition [34]. To avoid this problem, partial enumeration methods have been developed [34–36]. These methods explore the solution space (*i.e.*, the range of possible optimal condition-specific metabolic networks for this condition) to find a representative set of possible condition-specific metabolic networks. Hence, considering thousands of condition-specific metabolic networks is more robust and less error-prone than considering only one network in the solution space [34]. Nevertheless, it complexifies the interpretation since a condition will be associated with thousands of potential network configurations. There is thus a need to define a new strategy to extract mechanistic information from these numerous condition-specific metabolic networks to allow the final identification of the mMoA.

In this study, we propose a new semi-automated workflow designed to use transcriptomic data and GSMNs to identify and visualize mMoA. We used gene expression data from the Open TG-GATEs database to build condition-specific metabolic networks with a partial enumeration method adapted from the DEXOM [34] approach. Open TG-GATEs has been created from two projects spanning several years: Toxicogenomics Project One [37] (TGP1: 2004 to 2007)

and Toxicogenomics Project Two [38] (TGP2: 2010 to 2011). Open TG-GATEs contains *in vivo* gene expression and physiological data generated on Crl:CD Sprague–Dawley rats and *in vitro* gene expression data generated on PHH and primary rat hepatocytes. We selected gene expression data generated on primary human hepatocytes exposed for 24h to the maximum dose corresponding to a dose level yielding an 80-90% survival ratio. To compare two conditions, we developed a statistical analysis consisting of identifying DARs between two distinct sets of condition-specific metabolic networks representing the two conditions of interest and we took advantage of the GSMN structure to compute the metabolic distance between DARs in order to identify clusters of functionally interdependent metabolic reactions. We then extracted a minimal subnetwork for each cluster to better understand how these small biological processes are constructed and how reactions are interconnected. Finally, we used MetExploreViz [39] to interactively visualize the subnetworks and overlay additional annotation on the network such as cellular compartments, metabolic pathways, and custom mappings. This approach enables the analysis of the metabolic impact of a chemical at a global level but also at a very precise level such as the biochemical reaction scale. The workflow has been applied to eight molecules (ethanol, valproic acid, indomethacin, amiodarone, allopurinol, rifampicin, sulindac, and tetracycline) selected for their known hepatotoxic effects and the large amount of scientific literature describing their MoA. We further highlight these new computational workflow capabilities by predicting and analyzing the metabolic impact of two known hepatotoxic compounds: amiodarone and valproic acid.

2 Results

In this study we propose a new workflow to improve the understanding of the mMoA of chemicals. The workflow can be divided into three main steps.

The first step (Fig 1, box A) consists of building condition-specific metabolic networks. Indeed, GSMNs encompass all possible metabolic reactions regardless of the tissue, cell, or condition. Performing computational analysis on this generic model may raise inaccurate results (*e.g.*, highlighting a pathway which is known to be inactive for a cell type). Condition-specific metabolic networks can be built to overcome this limitation. They are composed of reactions predicted as active by a modeling algorithm for a given biological condition based on the generic GSMN and a set of gene expression data. In this study, we used Recon2.2 [26] as the initial GSMN to build reconstructions for several biological conditions (*i.e.*, cells exposed to a chemical) and transcriptomic data obtained from the Open TG-GATEs database (described in "Methods" in the section "Transcriptomic data processing"). For each studied condition, the method computes a set of condition-specific metabolic networks optimally matching gene expression data, network topology and the stoichiometry of reactions with an adapted version of DEXOM [34] (described in "Methods" in the section "Condition-specific modeling with partial enumeration"). DEXOM output is a dataset that gathers, for each reaction, the number of times it has been predicted as active by the algorithm. This result highlights the likelihood of each reaction of being active for a specific condition. Note that this step allows the prediction of activity for reactions, even if they are not associated with any gene (*e.g.*, passive transport reactions), by inferring their activity from the activity of surrounding reactions.

In the second step (Fig 1, box B) of the workflow, the objective is to identify reactions that activity changes in a significant manner between control and treated conditions. We therefore apply a statistical test between sets of reactions' predicted activity in order to identify significantly differentially activated reactions (DAR) between two conditions (described in "Methods" in the section "Identification of differentially activated reactions"). Since some reactions cannot be correctly constrained by transcriptomic data (*e.g.*, associated with less specific genes or no gene at all), they are more likely to be indiscriminately predicted as active or inactive while maintaining the same optimal fit with the data. The predicted activity of these reactions may therefore be more variable even across optimal subnetworks representing the same condition. We implemented a "baseline noise" calculation and filtration approach (described in "Methods" in the section "Baseline noise calculation") to remove these reactions from the

list of DARs obtained at the end of this step and to ensure their robustness. The list of DARs by itself is a very insightful result since it allows the listing of reactions whose activity is significantly affected by the studied condition (exposed *vs* control cells).

The last step (Fig 1, box C) of the workflow consists of deciphering where and how cell metabolism is perturbed by chemical exposure. DARs are not independent and may act in a coordinated manner through sequences of reactions, hence highlighting the potential modulated cascade of enzymatic reactions. To do so, a network-based approach has been developed to detect, visualize, and analyze the functional role of previously identified DARs. The method aims at stratifying the list of DARs into clusters of close reactions in the metabolic network (detailed in "Methods" in the section "DARs clustering"). Two reactions are considered to be close when they are connected to each other with only a few intermediary reactions. Once these clusters were identified, close reactions were used to extract small human-readable subnetworks (detailed in "Methods" in the section "Subnetwork extraction") that would describe parts of the metabolic network that are specifically modulated in the studied condition, suggesting potential mMoA of the molecules.

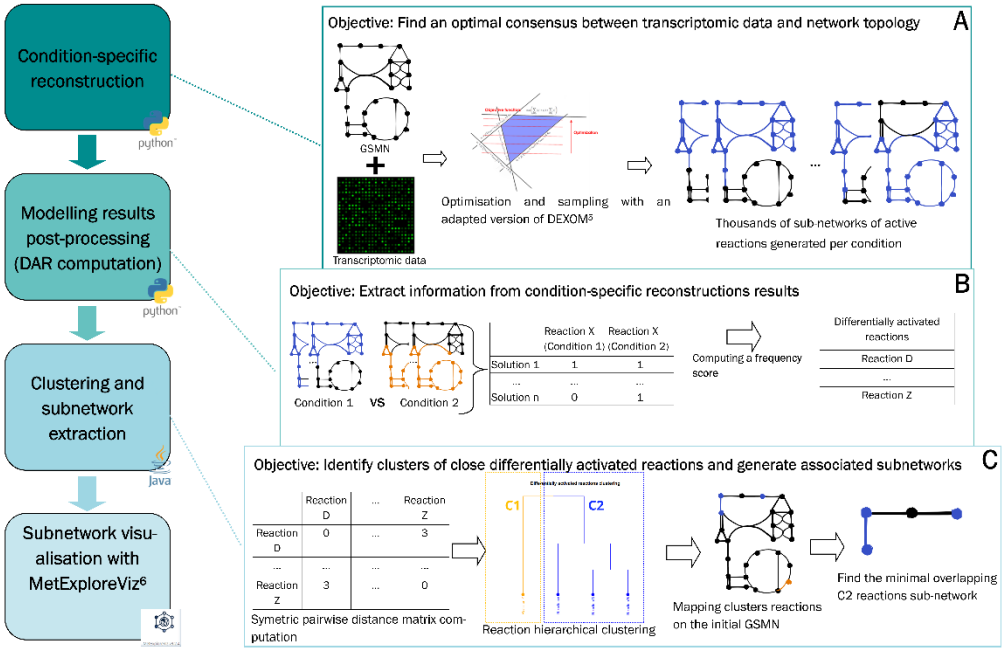

Fig 1. General overview of the three-step workflow. In the first step (A), transcriptomics data are integrated to a GSMN with a sampling approach adapted from DEXOM. In the second step (B), DARs are computed from the large number of metabolic networks obtained after the optimization and sampling step. Finally, in the third step (C), network analysis methods are developed and employed to interpret DARs and improve our understanding of the chemicals' mMoA.

Identification of differentially activated reactions (DAR) associated with exposure to eight chemicals

The workflow has been applied to eight molecules (ethanol, valproic acid, indomethacin, amiodarone, allopurinol, rifampicin, sulindac, and tetracycline) selected for their known liver toxicity and the associated MoAs reported in the literature. Gene expression data for the eight molecules from the Open TG-GATEs database were used. Data generated in PHH at the highest concentration after 24h of exposure for each compound were selected along with their associated controls. Condition-specific metabolic networks were reconstructed (Fig 1, box A) from gene expression data for each retrieved sample (two replicates per condition). On average, 10,000 different solutions

were enumerated for each sample. Each optimal solution is composed of a unique set of active and inactive reactions. In Table 1, the minimal, maximal, and average number of active reactions for the solutions calculated by DEXOM for each chemical and two vehicle controls (*i.e.*, medium and DMSO) are reported. It is worth noting that the size (*i.e.*, the number of active reactions) of DEXOM optimal solutions is in the same order of magnitude for all conditions (treatment or controls, see Table 1) with a minimal number of active reactions ranging from 3423 (DMSO condition) to 3639 (tetracycline 25µM, 24h), a maximal number of active reactions ranging from 4396 (valproic acid 5000µM, 24h) to 4645 (indomethacin 200µM, 24h), and a mean number of active reactions ranging from 4070 (valproic acid 5000µM, 24h) to 4570 (amiodarone 7µM, 24h).

	ethanol (10000µM, 24h)	valproic acid (5000µM, 24h)	indomethacin (200µM, 24h)	amiodarone (7µM, 24h)	allopurinol (140µM, 24h)	rifampicin (70µM, 24h)	sulindac (3000µM, 24h)	tetracycline (25µM, 24h)	medium	DMSO
Minimal number of active reactions	3475	3522	3498	3623	3534	3535	3499	3639	3578	3423
Maximal number of active reactions	4615	4396	4645	4570	4563	4620	4479	4580	4626	4610
Mean number of active reactions	4260	4070	4266	4570	4176	4284	4122	4209	4216	4192

Table 1. Description of the number of predicted active reactions for each chemical and vehicle set of condition-specific metabolic networks. *Sets of condition-specific metabolic networks were calculated by the adapted DEXOM approach implemented in our workflow (Fig 1, box A). All solutions obtained for samples of the same molecule have been merged before calculating the minimal, maximal, and mean number of active reactions per condition.*

DARs were computed to identify metabolic reactions modulated following exposure to each selected molecule (Fig 1, box B) compared to the corresponding control condition. The number of identified DARs ranged from 57 for PHH exposed to indomethacin (200µM, 24h) to a maximum of 477 for PHH exposed to valproic acid (5000µM, 24h) (Table 2). Only valproic acid, indomethacin, amiodarone and allopurinol DARs lists were slightly affected by the noise filtration procedure with 12.6%, 5.3%, 6.7% and 11.4%, respectively, of DARs filtered out (Table 2). Conversely, for tetracycline and ethanol, 43.4% and 62.8% of DARs were respectively filtered out (*i.e.*, reactions being modulated by both molecules and control conditions (S1 Table)). The DARs specificity ratio was also calculated for each condition that corresponds to the number of DARs retrieved solely in this condition and not in any other studied condition (*i.e*., DARs specific to the condition) divided by the total number of DARs retrieved for this condition (S2 Table). The DARs specificity ratio ranged from 22.2% for indomethacin to 91.1% for valproic acid. The DARs specificity ratio indicated that the predicted mMoA for indomethacin was at least similar to one other condition in the study; in contrast, the predicted mMoA for valproic acid was not similar to any other condition in the study.

	ethanol (10000µM, 24h)	valproic acid (5000µM, 24h)	indomethacin (200µM, 24h)	amiodarone (7µM, 24h)	allopurinol (140µM, 24h)	rifampicin (70µM, 24h)	sulindac (3000µM, 24h)	tetracycline (25µM, 24h)
Number of DARs	94	477	57	60	88	121	242	99
Number of DARs after filtration	35	417	54	56	78	98	181	56
% of reactions filtered out	62.8	12.6	5.3	6.7	11.4	19	25.2	43.4

Table 2. Number of DARs obtained for each condition. *DARs were identified by the developed workflow for PHH after 24h exposure to eight molecules known for their hepatotoxicity at dose levels yielding an 80-90% survival ratio.*

Global analysis of DARs associated with exposure to amiodarone and valproic acid

Two of the eight molecules were selected for further analysis: amiodarone (7µM, 24h) and valproic acid (5000 µM, 24h). These two molecules were selected because 1) they are well described hepatotoxic compounds with several MoAs known to induce liver damage such as steatosis [40,41], and 2) their level of transcriptomics-based information differs, with 5709 and two differentially expressed genes (DEGs) for PHH exposed to valproic acid and amiodarone, respectively (described in "Methods" in the section "DEG identification"). Despite the low number of DEGs observed for amiodarone, 56 DARs were identified in PHH after 24h exposure to 7µM amiodarone using the developed method.

Fig 2 shows the visualization produced by MetExploreViz for DARs identified for amiodarone and valproic acid in the context of human GSMNs using Recon 2.2. For amiodarone (Fig 2 A), a main group is composed of well interconnected up-activated reactions, while few other groups of up- or down-activated DARs are disconnected with no shared substrates or products. For valproic acid (Fig 2 B), many small groups of DARs are disconnected from each other without any particular area of the metabolic network specifically impacted. Hence, this general observation indicates that the mMoAs differ between these two molecules, and further in-depth analysis is required.

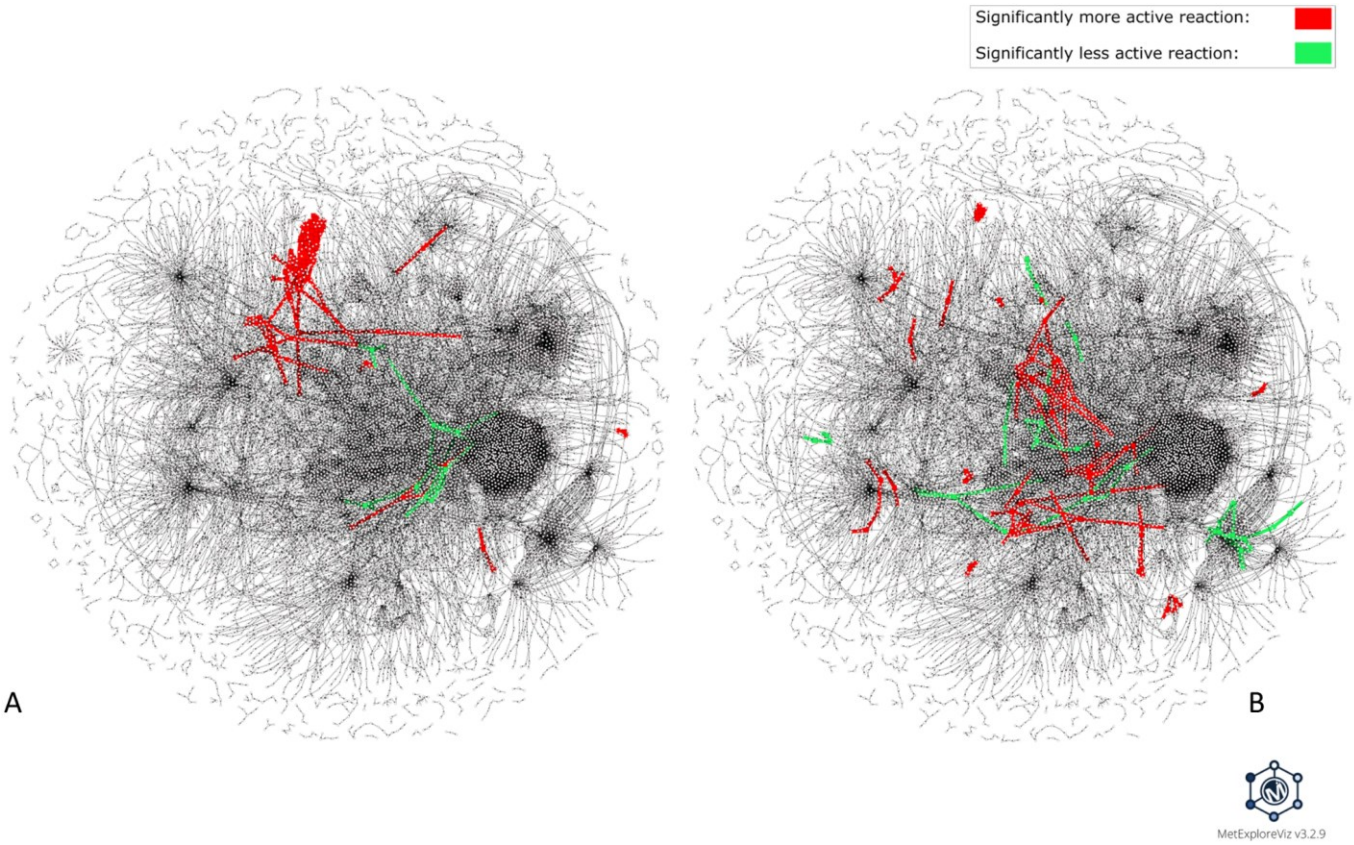

Fig 2. Visualization of DARs identified for amiodarone and valproic acid within the Recon2.2 metabolic network. This visualization was performed with MetExploreViz while removing side compounds (S3 Table). DARs identified for amiodarone (7µM, 24h) are highlighted in Fig 2 A and DARs identified for valproic acid (5000µM, 24h) are highlighted in Fig 2 B, with more frequently active reactions in the exposed *vs* control condition colored in red and less frequently active reactions in green. Nodes represent reactions and metabolites and are connected if a metabolite is a product or a substrate of a reaction. Fig 2 A and Fig 2 B are both based on the same network layout; thus, each reaction and metabolite is located at the same coordinates, allowing visual comparison.

Identifying subnetworks of DARs with graph-distance clustering and subnetwork extraction

In order to capture the mMoA of a molecule without having to rely on subjectively defined pathways [11,12], it is assumed that the proximity within the network can be used as a measure of the metabolic interdependency between reactions. Indeed, reactions are linked through compounds being produced and consumed by others. The shorter the chain between two reactions, the stronger the expected interdependency. Therefore, the metabolic distance, which is the length of the chain or path between two reactions in the metabolic network, is used to estimate interdependency. Hence, we propose grouping reactions that are closely located as a proxy to identify reactions involved in the same metabolic function.

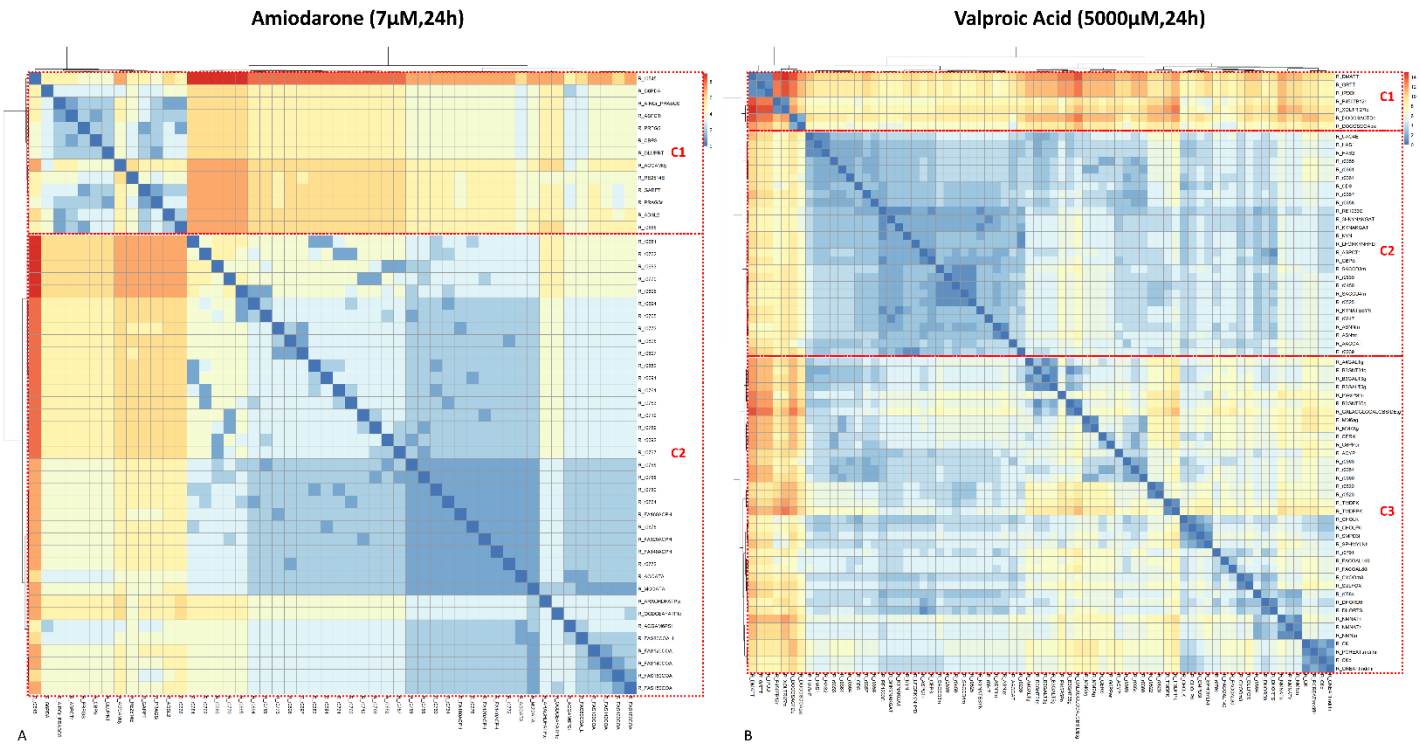

Fig 3. Biclustered heatmap of the pairwise reaction distance matrix for amiodarone and valproic acid. Hierarchical biclustering on the pairwise reaction distance matrix for amiodarone and valproic acid was performed with the Ward algorithm. The resulting biclustering was visualized as a heatmap computed with the Pheatmap R package. The color scale depicts the distance between two reactions. The distance ranges between zero (cells colored in blue) to eight for amiodarone (Fig 3 A) and 14 for valproic acid (Fig 3 B) (cells colored in red). Two main clusters (C1 and C2) were identified for amiodarone (Fig 3 A) and three clusters (C1, C2, and C3) were identified for valproic acid.

The pairwise distance matrix between all reactions identified as differentially active for each of the two conditions, amiodarone and valproic acid, was computed based on the list of DARs. As described in the "Methods" section, the computation of pairwise distance was performed on the metabolic reaction undirected graph of Recon2.2, using the shortest path distance metric. Then, subsets of DARs were identified by clustering the pairwise distance matrix with a hierarchical clustering approach. Two clusters were selected for PHH exposed to 7µM amiodarone for 24h (Fig 3 A) and three clusters for PHH exposed to 5000µM valproic acid for 24h (Fig 3 B).

To go further in the interpretation and to better understand how these reactions interact and are implicated in the mMoA of our chemicals of interest, a subnetwork extraction step was implemented and applied to each cluster of DARs. Extracting a subnetwork allows one to visualize how the DARs are interconnected and fill gaps between disconnected DARs by adding the necessary intermediate reactions. The subnetwork extraction was performed using the Met4J library (https://forgemia.inra.fr/metexplore/met4j) implementation of the minimal Steiner tree on a pruned reaction graph of Recon2.2. The minimal Steiner tree extraction algorithm enables the extraction of a subnetwork connecting two lists of nodes (*i.e.*, DARs) while minimizing the size of the final subnetwork. Therefore, this algorithm well fits our objective of visualizing interactions between DARs while keeping it human-readable. To guide our analysis, the "DARs subnetwork coverage" metric was defined, which is calculated as the number of DARs in a subnetwork divided by the total number of reactions in this subnetwork. This metric provides an estimate of how rich in DARs a subnetwork is. A high DARs subnetwork coverage score indicates closely located and tightly interacting DARs.

For PHH exposed to 5000µM of valproic acid for 24 hours, a subnetwork with a 77% DARs subnetwork coverage was extracted (Cluster C2 in Fig 3 B, see S4 Table for details), meaning that most reactions within the extracted

subnetwork were DARs that were closely interconnected. Among the DARs of this subnetwork, 21 were up-activated while only four were down-activated (Fig 4 A). Five DARs predicted as up-activated were associated with -oses metabolism such as fructose/mannose metabolism and glycolysis/gluconeogenesis pathway, some of which are involved in the phosphorylation of hexoses (Fig 4 B). Another group of up-activated DARs is associated with lysine metabolism in mitochondria (html links for visualization in S1 Appendix) and especially the degradation of lysine through production of L-saccharopinate from 2-oxoglutarate. One up-activated reaction is associated with transport of L-asparagine into the mitochondria, which is then used by the L-asparaginase to produce L-aspartate. The resulting L-aspartate is finally transported by an aspartate-glutamate shuttle that was not differentially modulated. The phosphatidate cytidylyltransferase associated with the glycerophospholipid metabolism was also predicted to be up-activated in PHH after exposure to 5000 µM valproic acid for 24 hours. Finally, eight differentially up-activated reactions were associated with tryptophan metabolism and involved in reactions associated with kynurenate, L-glutamate and 2-oxoglutarate. Down-activated reactions were more sparsely located in the extracted subnetwork and were associated with hyaluronan metabolism, propanoate metabolism, where acetoacetate and coenzyme A are conjugated to produce acetoacetyl-CoA, and finally with pyrimidine synthesis, where glutamate and aspartate production/consumption were disturbed. The subnetwork extraction algorithm could add reactions that were not identified as DARs but were necessary for the connectivity of the subnetwork. These added reactions were associated with pathways also associated with DARs such as mitochondrial transport, glycolysis/gluconeogenesis, pyrimidine synthesis and lysine metabolism (Fig 4 B).

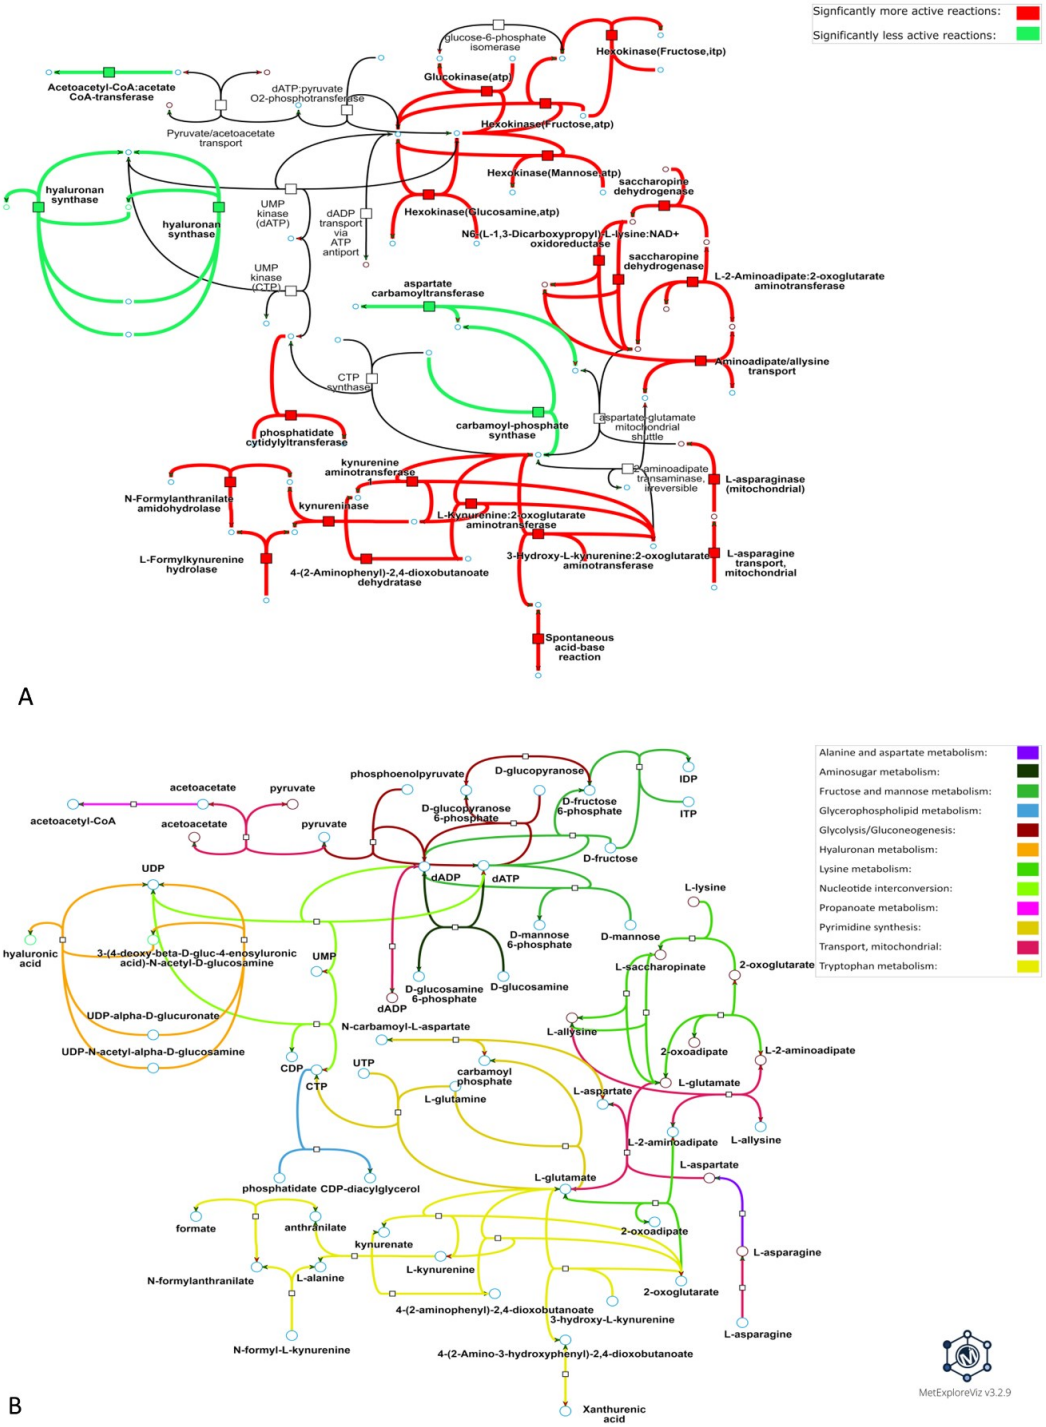

Fig 4. Metabolic visualization of the subnetwork extracted from a cluster of DARs obtained from the condition-specific reconstruction process in PHH exposed to 5000µM valproic acid for 24h. The visualized subnetwork was computed from DARs in cluster 2 of valproic acid (5000µM, 24h), which is the cluster with the highest DARs subnetwork coverage among the clusters identified in the distance matrix (see Fig 3 B) for this condition. Nodes represented by a square are metabolic reactions and nodes represented by a circle are metabolites. Fig 4 A and Fig 4 B represent the same subnetwork with the same topology. Links in Fig 4 A are highlighted according to the direction of the perturbation (*e.g.*, if the reaction is more frequently active in the exposed *vs* control condition) and links in Fig 4 B are colored according to the metabolic pathway of the associated reaction (details in legend). Interactive visualizations can be accessed through these links:

https://metexplore.toulouse.inrae.fr/userFiles/metExploreViz/index.html?dir=/72ff7fdc7031b880ef4f3532134aa326/networkSaved_373423088
https://metexplore.toulouse.inrae.fr/userFiles/metExploreViz/index.html?dir=/72ff7fdc7031b880ef4f3532134aa326/networkSaved_725935955

For PHH exposed to 7µM of amiodarone for 24 hours, a subnetwork with a DARs coverage of 95% (S4 Table) was extracted, indicating that nearly all the reactions represented in Fig 5 are DARs. In Fig 5 A, it was observed that 36 DARs were up-activated in the treated condition and only one DAR was down-activated in the treated condition. Most of the up-activated reactions (*i.e.*, 31) of this subnetwork (Fig 5 A) are associated with the fatty acid synthesis, four reactions are linked to the fatty acid oxidation pathway and one reaction is associated with aminosugar metabolism. Regarding the fatty acid synthesis pathway, many reactions associated with conjugation/deconjugation of acyl carrier protein (ACP) to fatty acids were perturbed. Fatty-acyl-CoA synthase reactions were also perturbed with many reactions involved in the addition of malonyl-CoA to fatty acids as being differentially up-activated. Some up-activated reactions associated with the fatty-acid oxidation pathway were also disturbed, such as the acetyl-CoA-ACP transacylase and the malonyl-CoA-ACP transacylase as well as two fatty-acyl-ACP hydrolases. Finally, the only down-activated reaction of this subnetwork is a reaction associated with the aminosugar metabolism involved in N-acetylglucosamine-6-phosphate's production from acetyl-CoA and D-glucosamine-6-phosphate. In the subnetwork presented in Fig 5, only two reactions associated with fatty-acyl-CoA elongation from the fatty-acid synthesis pathway were added by the subnetwork extraction algorithm.

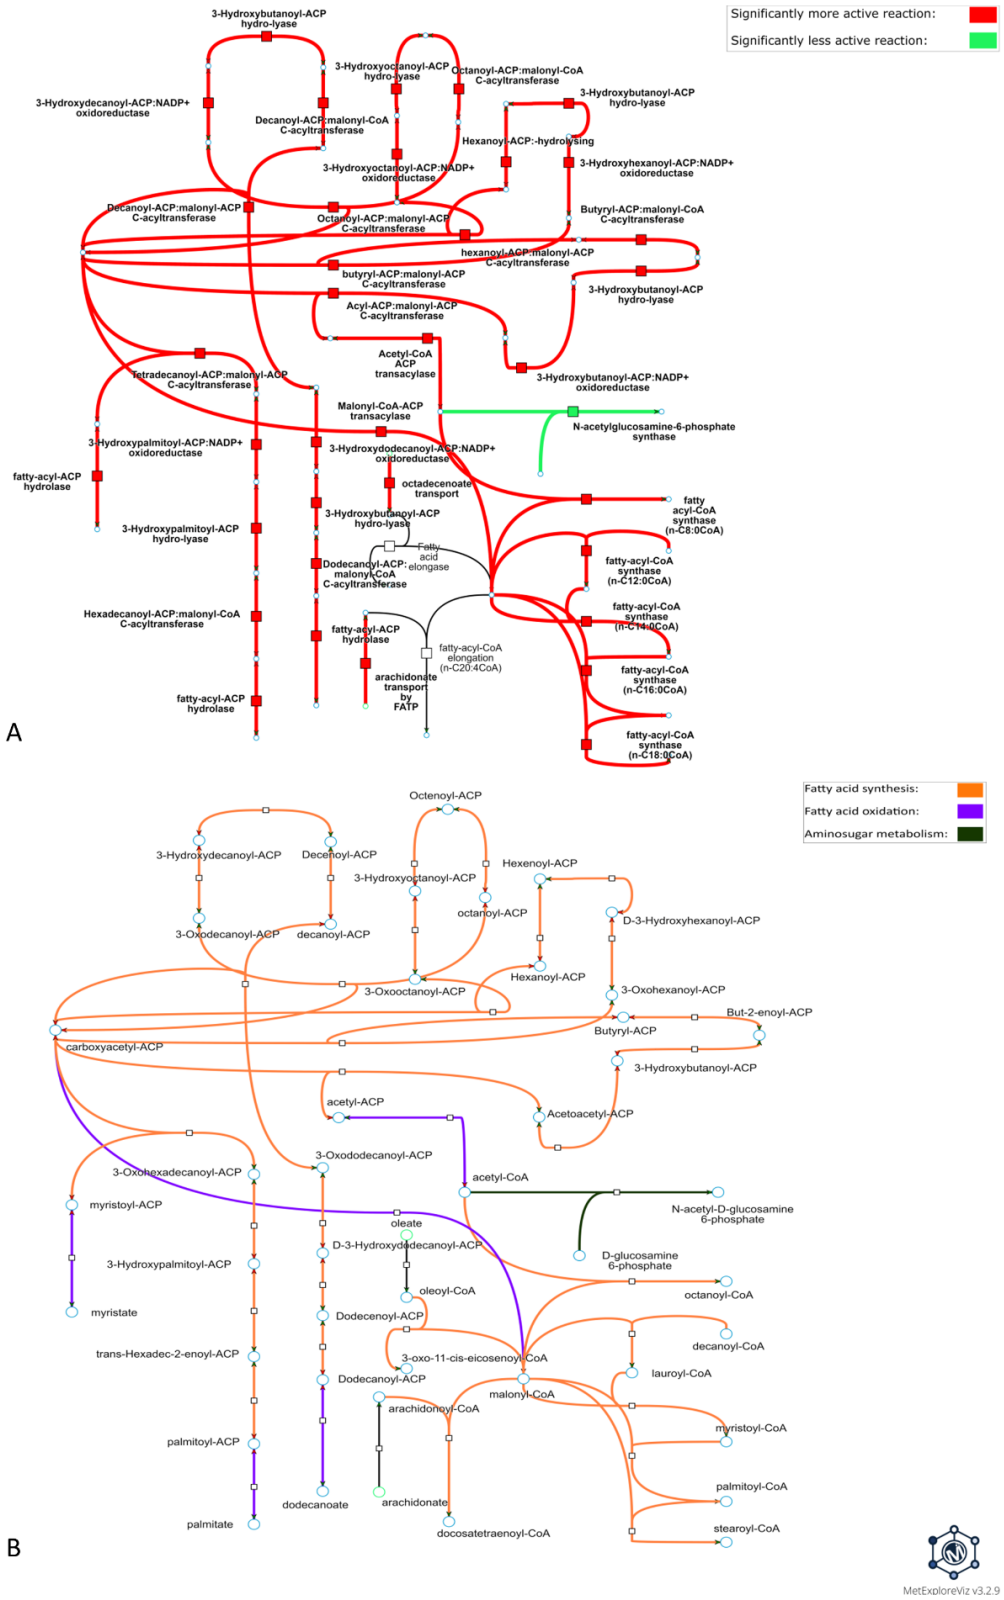

Fig 5. Metabolic visualization of the subnetwork extracted from a cluster of DARs obtained from the condition-specific reconstruction process in PHH exposed to 7µM amiodarone for 24h. The visualized subnetwork was computed with DARs from cluster 2 of amiodarone (7µM, 24h), which is the cluster with the highest DARs subnetwork coverage among the clusters identified in the distance matrix (Fig 3 A) for this condition. Nodes represented by a square are metabolic reactions and nodes represented by a circle are metabolites. Fig 5 A and Fig 5 B represent the same subnetwork with the same topology. Links in Fig 5 A are highlighted

according to the direction of the perturbation (*e.g.*, if the reaction is more frequently active in the exposed condition *vs* control condition) and links in Fig 5 B are colored according to the metabolic pathway of the associated reaction (details in legend). Interactive visualizations can be accessed through these links:
https://metexplore.toulouse.inrae.fr/userFiles/metExploreViz/index.html?dir=/72ff7fdc7031b880ef4f3532134aa326/networkSaved_292937465
https://metexplore.toulouse.inrae.fr/userFiles/metExploreViz/index.html?dir=/72ff7fdc7031b880ef4f3532134aa326/networkSaved_1994092833

3 Discussion

Calculation of DARs from the sampling of condition-specific reconstructions

To characterize and visualize the mMoA of chemicals, we developed a workflow based on condition-specific reconstructions and network analysis. Condition-specific reconstructions integrate gene expression data to GSMNs that are representative of the general endogenous metabolism of an organism. Such an integration allows additional information to be uncovered that could not be retrieved from transcriptomics data, such as the activity of reactions not associated with the expression of specific genes (passive transports, spontaneous reactions, etc.) and cellular localization. To be predicted as active, a reaction must have its source metabolite(s) being produced by another reaction in the metabolic network, and its product(s) being consumed by another reaction. Therefore, the objective of the modeling algorithm is to find the best consensus between reactions predicted as active according to gene expression data and these production/consumption constraints from the metabolic network. However, as mentioned previously, the amount of information available in gene expression data for the modeling is not sufficient to allow the algorithm to find a unique condition-specific network. This limit has already been described by several authors [34–36,42] and partial enumeration approaches have been developed to circumvent this issue. Such methods provide many similarly adequate optimal subnetworks based on the gene expression data, which raises the question of how to analyze the different optimal solutions obtained. Rossell *et al.* proposed EXAMO to reconstruct a minimal condition-specific network from the alternative optimal solution space that prioritizes reactions found to be active in all alternative solutions. This strategy tends, however, to emphasize central metabolism although it might not necessarily be the part of metabolism most affected by chemical compound exposure. Poupin *et al.* identified required, inactive and potentially active reactions, based on how frequently the reactions were found to be active in the alternative optimal solution space. In this method, the group of reactions defined as potentially active is difficult to interpret and ended up being merged with required reactions to perform the pathway enrichments analyses in their study. To consider reaction activities more globally, we introduce the DAR calculation approach, which aims to compare two sets of condition-specific metabolic networks and identify reactions that are significantly perturbed in one condition compared to another (*e.g.*, treated *vs* control). However, as described earlier, some metabolic reactions are more sensitive to the under-constrained problem and will be predicted as active or inactive independently of experimental data specificity. To limit this effect, we introduced a filtration step called “baseline noise filter” that takes into account the uncertainty of predicting the activation status associated with some reactions. The baseline noise filter had a different impact on DARs lists, depending on the condition. PHH exposed to ethanol was the most impacted condition with 62.8% of its DARs removed by the filter, while PHH exposed to indomethacin was the least impacted condition with 5.3% of its DARs removed. Most of the DARs filtered out for ethanol and tetracycline (*i.e.*, the most impacted conditions) are transport and peroxisomal β-oxidation reactions. Indeed, these reactions are more sensitive to the under-constrained problem either because they are not associated with any genes (*i.e.*, exchange reactions) or they are associated with complex GPR rules that tend to result in unconstrained reactions. The baseline noise filter is a rather conservative approach, but it ensures that the predicted perturbation is not due to a lack of constraints on the model. Indeed, one key point in the development

of condition-specific reconstructions was to get enough biological constraints. In this study, we used transcriptomic data only, but our reconstruction process could be further refined by combining transcriptomic data with medium metabolomics data (exometabolomics) that would add further constraints on metabolic fluxes, thereby reducing the uncertainty related to the number of possible alternative solutions.

Comparison and complementarity with pathway enrichment

To evaluate the added value of our workflow, in comparison to directly using gene expression data through approaches such as pathway over-representation on DEG, we performed a pathway over-representation analysis on DEGs for the valproic acid condition (S1 Fig). The first drawback to the DEG-based approach is that it becomes inapplicable or unreliable when the list of DEGs is short, which is often the case for cells exposed to a low concentration of chemicals triggering subtle metabolic impacts, such as 24h treatment of PHH with 7µM amiodarone in our study with only two DEGs. On the contrary, we showed that our workflow is able to gain information from transcriptomic data even when the DEGs list is short. Nevertheless, when performing pathway over-representation on DEGs with Reactome (*i.e.*, including all genes, not just metabolic genes) for PHH exposed to 5000µM valproic acid for 24h, we observed that many of the top 50 enriched terms are related to cell cycle, gene expression regulation, and cell signaling (S1 Fig). Although this is valuable information about the MoA of valproic acid, it is not directly associated with its metabolic impact in liver cells, suggesting that combining both gene expression data analysis and condition-specific reconstruction is beneficial. We performed a pathway enrichment study with DARs on Recon2.2 pathways (S2 Fig) for the eight molecules: we observed that the fatty acid synthesis pathway was evidenced as significantly enriched from DARs for four molecules and compared how these four molecules were impacting this pathway (Fig 6) by visually comparing the metabolic footprint of each chemical on the fatty acid synthesis pathway. Amiodarone and allopurinol affected the same reactions (Fig 6 A and Fig 6 C), which suggested that they share similar mMoAs regarding fatty acid synthesis. On the other hand, rifampicin and sulindac seemed to impact another part of the fatty acid synthesis pathway (Fig 6 B and Fig 6 D). These observations showed that, although pathway enrichment can provide information on which general pathways are impacted by a chemical, it is not sufficient to evidence more precise modulated functions or to differentiate between potential different mechanisms of action. Even if these four compounds are identified as being able to induce liver steatosis, underlying MoAs might differ [43]. It is then interesting to see how our approach might allow in-depth understanding of MoAs.

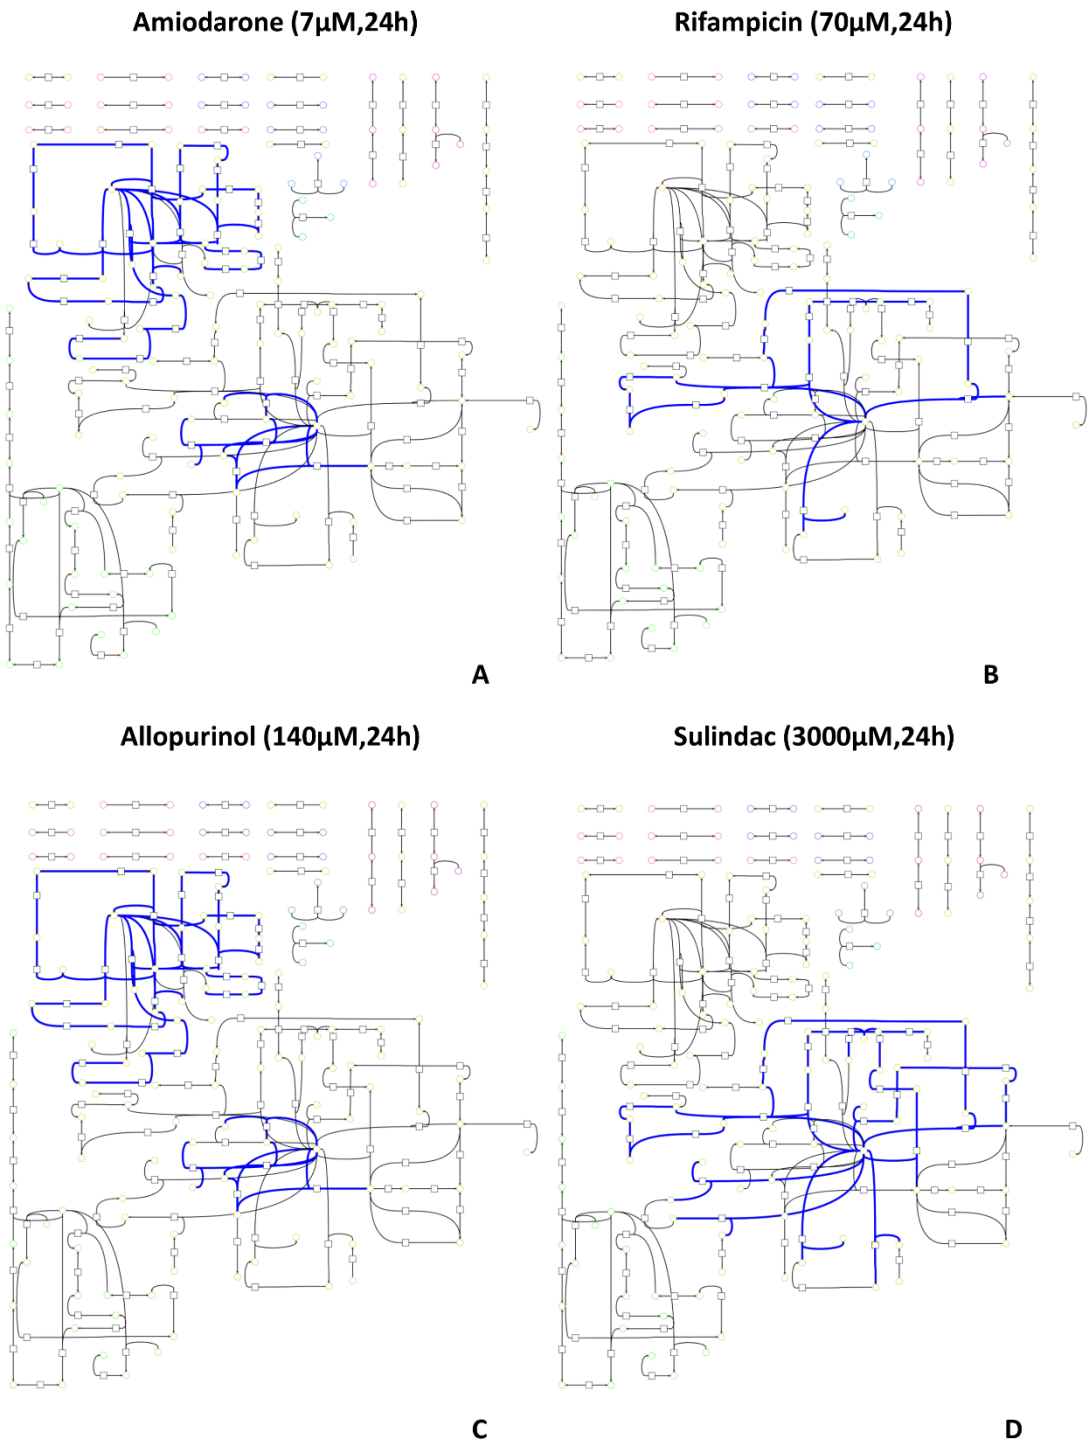

Fig 6. Comparison of the metabolic impact on the fatty acid synthesis pathway for four molecules for which this pathway is significantly enriched. Each metabolic graph represents which fatty acid pathway reactions are differentially activated (in blue) after *in vitro* exposure to amiodarone (Fig 6 A), sulindac (Fig 6 B), allopurinol (Fig 6 C), and rifampicin (Fig 6 D). Nodes represented by a square are metabolic reactions and nodes represented by a circle are metabolites. DARs links have been highlighted in blue to identify which part of the pathway is perturbed by the molecule.

Metabolic distance-based clustering and subnetwork extraction

To go beyond pathway enrichment-based analysis, we propose a clustering approach based on the metabolic distance in the network, which ensures the identification of groups of functionally interdependent modulated

reactions. These groups can involve reactions belonging to different pathways, therefore enabling the study of the mMoA as a continuous phenomenon instead of a binary (*i.e.*, enriched/not enriched pathway) phenomenon. However, hierarchical clustering suffers from some limitations. Indeed, hierarchical clustering is not designed to capture communities (*i.e.*, groups of densely interconnected elements in a graph) and tends to consider reactions implicated at the ends of linear cascade as distant, therefore not biochemically related, even if according to the network topology these reactions are biochemically related and highly dependent. Community detection approaches could be an answer to some of these limitations but their performance may be affected by the particular topological properties of large networks [44] such as GSMNs, therefore requiring a thorough evaluation to find the most effective method.

To add metabolic context and functional information to the lists of DARs and find connected subnetworks that include clustered DARs, we extracted minimal subnetworks with an algorithm based on the Steiner minimal tree problem. This algorithm is well adapted to our principal objective of visualizing human-readable subnetworks because it will find the smallest subnetwork connecting all the nodes of interest (*i.e.*, DARs) while adding as few nodes as possible for connectivity. However, the Steiner minimal tree problem is computationally hard (proven to be NP-Complete [45]). The metric closure graph [46] (or shortest distance graph) approximation was thus used. While employing an approximation is necessary, it introduces an arbitrary decision-making process within the algorithm, leading to the selection of a single alternative minimal Steiner tree among various possibilities. This choice, although necessary to extract a minimal subnetwork representative of the mMoA, could miss biologically interesting connections between DARs. The drawback of favoring the obtention of small graphs is that some interesting alternatives could be omitted. Another option could therefore be the k-shortest paths algorithm, which would make it possible to find these alternatives, at the expense of readability. Hence, regarding subnetwork extraction, method choice is guided by finding the right trade-off between readability and exhaustivity.

Because metabolic disruptions can impact a wide range of cellular functions spanning several metabolic pathways, the clustering and subnetwork extraction methods are key to connect and analyze disruptions. Interestingly, in the subnetwork presented in Fig 4, most of the reactions added by the subnetwork extraction algorithm are transport reactions, which highlight how the subnetwork extraction algorithm can help to better understand the mMoA of the molecule as a continuous phenomenon spanning several pathways and compartments rather than localized and disconnected perturbations.

PHH exposed to amiodarone or valproic acid subnetworks mMoA analysis

Amiodarone and valproic acid are two well-studied drugs known to induce hepatotoxicity via the occurrence of steatosis [40,41], which is in line with perturbations predicted by our workflow and visualized in Fig 4 and Fig 5 associated with -oses and lipid metabolism. Valproic acid is also widely known for its impact on mitochondrial function [41] and is a histone de-acetylase inhibitor [47], a class of molecules known to impair many cell functions; this is supported by the wide diversity of pathways subparts identified in the mMoA of valproic acid (Fig 4B).

To study the metabolic impact of valproic acid on PHH after 24h exposure, we focused on visualizing cluster 2 (Fig 4), which is the cluster with the highest DARs coverage for this condition. We observed that DARs from this cluster were associated with 12 pathways. We were also able to identify another group of four up-activated reactions that is associated with lysine degradation in mitochondria, a phenomenon associated with mitochondrial homeostasis disruption in mice [48]. This phenomenon could also be due to a decrease in the L-carnitine pool associated with valproic acid exposure [49,50] that could trigger a compensating mechanism to restore L-carnitine which requires lysine as a substrate. Finally, a group of eight reactions predicted as up-activated is associated with the metabolism of tryptophan. Interestingly, an increase in the metabolism of tryptophan and kynurenine has already been reported as a potential effect of valproic acid in rats [51], and the valproic acid-induced increased conversion of tryptophan to nicotinamide has been observed in rats [52].

The impact of amiodarone on the fatty acid synthesis pathway predicted by our workflow (Fig 5) has already been observed in HepaRG cell cultures [43]. It is suggested that the increase in fatty acid synthesis could be due to the activation of SREBP1, which is a transcription factor mediating the expression of several important genes for *de novo* lipogenesis. One of these genes, FASN, encodes the fatty-acid synthase, which our workflow predicted as DAR in the treated state (Fig 5 A). Interestingly, SREBP1 and FASN were not identified as differentially expressed in amiodarone's transcriptomic data, suggesting that our workflow is able to retrieve mechanistic information from lowly informative transcriptomic data. A similar increase of *de novo* lipogenesis on 3T3L1 adipocytes has also been reported [53] and is linked to an increase in palmitate production, which is in line with the predicted up-activation of the fatty-acyl-ACP hydrolase hydrolyzing palmitoyl-ACP to palmitate (Fig 5). In the subnetwork presented in Fig 5, only two reactions have been added by the subnetwork extraction algorithm, suggesting that DARs identified in this cluster are closely located in the metabolic network, therefore functionally interdependent. Moreover, the visualized subnetwork accounts for 74% of predicted DARs for amiodarone. Since the majority of DARs in the visualized subnetwork (Fig 5) are associated with fatty acid synthesis, predictions from our workflow suggest that amiodarone mMoA is mostly related to this pathway and quite localized in the metabolic network. This mechanistic observation would not have been possible with the initial transcriptomic data analysis that yielded only two DEGs for PHH exposed 24h at 7µM amiodarone, thus emphasizing the capacity of our workflow to extract mechanistic information from condition-specific metabolic networks constructed with lowly informative transcriptomic data.

4 Conclusion

In this study, we developed a workflow that allows a better understanding of a chemical's mMoA by using metabolic gene expression data combined with several metabolic network modeling methods. We integrated metabolic gene expression data to a human GSMN to construct sets of condition-specific metabolic networks that we then analyzed through an original statistical approach to identify reactions that are likely to be differentially activated between the two compared conditions. To go further on the mechanistic understanding, we developed a graph-based network analysis approach to extract and visualize subnetworks corresponding to metabolic functional and connected groups of DARs. Finally, we showed that the strategy that we developed succeeded in retrieving parts of the mMoA described in the literature for PHH exposed to two well-known hepatotoxic molecules (amiodarone and valproic acid), even when the amount of available information in transcriptomic data was limited.

The workflow presented in this study allows a better and deeper understanding of chemicals' mMoA. This precise metabolic network exploration and visualization comes at the expense of the computation time, which could make it challenging to apply in large-scale safety assessment studies and would require development in parallel of complementary methods providing a more global assessment of the metabolic impact of a compound. These complementary developments could be plugged into the current workflow quite easily thanks to the flexibility of python and jupyter notebooks. Although we focused on studying the metabolic impact of chemicals on PHH, our workflow could also be used to model the metabolic impact of chemicals on different cellular types from different tissues, therefore paving the way for a more precise understanding of how a chemical impacts the metabolism of many organs.

5 Methods

Transcriptomic data processing

Transcriptomic data used in this study were obtained from the Open TG-GATEs [21] database. For this study, we selected *in vitro* data generated on PHH. Raw transcriptomic data were downloaded as CEL files from https://dbarchive.biosciencedbc.jp/en/open-tggates/download.html. CEL files were read with the affy [54] package and the resulting dataset was normalized with the robust multiarray average method [55] from the limma R package. Information about the different PHH cell batches used in the different experiments was obtained from the authors

(S5 Table) and was used to correct the batch effect with Combat [56], available in the SVA package. We annotated the probes with the annotation database corresponding to the Affymetrix HG133U Plus 2 chips available in the hgu133plus2.db package [57] and the AnnotationDbi package [58]. Where several different probes mapped to the same gene, we selected the probe with the highest standard deviation of expression values in the dataset. Finally, to match DEXOM [34] requirements, we categorized transcriptomic data with Barcode [59–61]. We used the z-scores output available from the barcode package and applied a 25/75 percentiles cutoff to obtain a list of highly and lowly expressed genes. Genes below the 25th percentile were considered as lowly expressed genes and assigned a -1 value, while genes above the 75th percentile were considered as highly expressed genes and assigned a +1 value. Genes between the 25th and 75th percentiles were not considered (*i.e.*, 0 value) and therefore did not have any impact on the modeling process.

DEG identification

Lists of DEGs were obtained from the ToxicoDB [20] database, which provides pre-processed differential gene expression data from several databases, including Open TG-GATEs. A jupyter notebook querying ToxicoDB API was developed to automate the process. First, it downloads the ToxicoDB compounds json file and iterates over the list of compounds obtained from this file to query the ToxicoDB API to get analyzed data associated with each compound. The retrieved data were then filtered by selecting genes from the Open TG-GATEs human study with an absolute value of log2(FC) larger than 0.26 and a q-value less than 0.05, from samples subjected to a "high" dose (three doses were investigated and available for this dataset: low, middle, and high) over 24 hours. The filtered DEGs list for each compound was then stored in a .tsv file.

Metabolic model preparation

We used the Recon2.2 [26] (downloaded from https://www.ebi.ac.uk/biomodels/MODEL1603150001) GSMN as the initial framework for the modeling and network analysis steps. This GSMN is composed of 5324 metabolites, 7785 reactions, and 1675 genes. We modified the model biomass reaction, to account only for the precursors required for cell maintenance but not replication, in order to better represent the fact that PHH are differentiated cells with a short lifespan and do not proliferate under classical culture conditions [62]. To do so, we set the stoichiometric coefficient of the DNA precursor in the biomass reaction to zero and applied a correction coefficient (S1 Appendix) to other metabolites to keep the reaction balanced. Next, we forced the lower bound of the modified biomass reaction to a value of 1 instead of 0 to ensure that the generated models will be able to produce biomass (*i.e.*, have a non-zero flux through the biomass reaction).

Condition-specific modeling with partial enumeration

We integrated the categorized transcriptomic data into the GSMN through GPRs to define an *a priori* set of active and inactive reactions. A highly expressed gene will be given the value of 1 whereas a lowly expressed gene will be given the value of -1. Genes are associated with reactions by the GPRs. GPRs are Boolean rules that indicate which genes are required to produce a specific enzyme that catalyzes one or more reactions. When a reaction has an AND GPR, the algorithm will annotate the reaction as active if the minimal value of the categorized GPR's gene expression values equal 1 (*i.e.*, all genes of the GPR are highly expressed) and inactive otherwise, such as:

$For\ Gene1\ =\ 1, Gene2\ =\ -1, Gene3\ =\ 1, Gene4\ =\ -1$:

$$Gene1\ AND\ Gene2 = min(Gene1, Gene2)\ =\ min(1, -1)\ =\ -1$$

In this case, if *Gene1* is highly expressed *AND Gene2* is lowly expressed, the reaction is considered inactive.

When a reaction has an OR GPR, the algorithm will annotate the reaction as active if the maximal value of the categorized GPR's gene expression values equal one (*i.e.*, at least one gene of the GPR is highly expressed) and inactive otherwise, such as:

$$Gene1\ OR\ Gene2 =\ max(Gene1, Gene2)\ =\ max(Gene1, Gene2)\ =\ max(1, -1)\ =\ 1$$

In this case, if *Gene1* or *Gene2* is highly expressed, the reaction is considered active.

These rules can be applied to more complex GPRs such as:

$$\begin{aligned}\big((Gene1\ OR\ Gene2)AND\ (Gene3\ OR\ Gene4)\big) &= \min(\max(Gene1, Gene2), \max(Gene3, Gene4)) \\ &= \min(\max(1, -1), \max(1, -1)) = \min(1,1) = 1\end{aligned}$$

In this case, if *Gene1* or *Gene2* is highly expressed and *Gene3* or *Gene4* is highly expressed, the reaction is considered active.

A reaction will be identified as *a priori* active if the resulting value equals 1 and as *a priori* inactive if the resulting value equals -1. At the end of this process, we identified a list of *a priori* active/inactive reactions for each sample that was then provided to DEXOM. DEXOM is a constraint-based modeling algorithm based on iMAT [29,30] whose objective is to find a steady state distribution of flux that maximizes the number of reactions whose flux is consistent with transcriptomic data levels [30].

To extract a representative set of solutions from the whole solution space, we adapted the partial enumeration from DEXOM and used the python implementation available in dexom-python (https://github.com/MetExplore/dexom-python). First, we applied a full Reaction-enum strategy (see [34] for the complete description). This method iterates over all the reactions of the network, blocking each of them successively and solving the resulting mixed integer linear problem. Next, we applied the DEXOM Diversity-enum strategy, starting with a random solution picked per range of Reaction-enum solutions as a starting point. Diversity-enum aims to find new solutions that are gradually more different from the starting solution, allowing one to explore the limits of the solutions' space. To reduce the computational costs of the original DEXOM partial enumeration approach, we reduced the number of starting solutions to a set of 10% of the solutions obtained with the Reaction-enum approach. We used an adapted version of systematic sampling (*i.e.*, one random solution picked per batch of solutions, S1 appendix for details) to generate a starting set representative of the complete set of Reaction-enum solutions. Finally, all the optimal solutions for a sample were grouped in a single tabulation-separated file. The list of parameters used for running the DEXOM algorithm can be found in S1 Appendix.

Identification of differentially activated reactions

For each reaction, the predicted activation status across all optimal condition-specific metabolic networks enumerated by the DEXOM strategy (Table 1) is stored in a numeric vector. Therefore, for each condition, the output of the DEXOM enumeration is a matrix of binary vectors where columns correspond to reactions and rows correspond to enumerated optimal solutions (Table 3).

DEXOM solution	**Reaction 1**	**Reaction 2**	**Reaction 3**	**…**	**Reaction N-1**	**Reaction N**
Solution 1	0	0	0	…	1	1
….	...	…	…	…	...	...
Solution N	1	1	0	…	0	1

Table 3. Example of DEXOM output for one condition: list of optimal sub-networks obtained from enumeration, stored as binary vectors

For each reaction and each condition, we can compute an activation frequency value *f*, as the number of solutions in which the reaction is predicted to be active, over the total number of solutions. To compare the activation frequency values of two conditions, we introduce a metric called R2.

$$R2 = \left(\frac{nActive_{ctrl}}{nTotal_{ctrl}} - \frac{nActive_{trt}}{nTotal_{trt}}\right)^2 = (f_{ctrl} - f_{trt})^2$$

Where $nActive$ is the number of times the reaction is predicted active and $nTotal$ is the total number of solutions in the condition (*i.e.*, $nActive + nInactive$). f_{ctrl} is the activation frequency of a reaction under the control condition and f_{trt} is the activation frequency of the same reaction under the treated condition.

We considered reactions as differentially activated if the R2 value was higher than 0.2 when comparing treated *vs* non-treated conditions.

Baseline noise calculation

The absence of gene expression information for some reactions, and therefore the absence of constraints on those reactions when using the DEXOM algorithm, leads to uncertainty (or "noise") in the reaction activation frequency prediction. Indeed, in the absence of transcriptomic data on reactions, DEXOM can predict these reactions as active or inactive without any impact on the optimality score. This uncertainty could lead to biased conclusions when comparing two conditions. To avoid considering reactions that are loosely constrained as DAR, we implemented a method that aims at estimating the baseline noise associated with each reaction in the network. The noise refers to the variation of the activation frequency of each reaction between pairs of control conditions, which is due to the fact that the predicted activity of the reaction does not impact the consistency with the data. Because several control conditions and chemical exposure times have been used in this dataset, we computed the baseline noise for each reaction between control conditions at a specified time and a specified vehicle. Replicates were pooled for each molecule. Practically, for each reaction, we computed the R2 between all pairs of control conditions for a given vehicle at a given exposure time and calculated the median of all R2s for each reaction. Therefore, we obtained for each reaction a baseline noise estimation that is the median of all pairwise comparisons between selected control conditions.

Then we filtered out DARs with an R2 between two conditions of interest (*e.g.*, control *vs* treatment) that was less than two times the noise estimate for this reaction.

Metabolic reaction graph construction

GSMNs, which are large and complex metabolic networks, contain hub nodes (*i.e.*, nodes connected to many other nodes). These hub nodes correspond mainly to cofactor compounds, *i.e.,* metabolites required in many biochemical reactions (ATP, ADP, NADPH, etc.) therefore creating connections between many reactions, although these connections are often not biologically meaningful. This property of GSMNs is a challenge [63] for network analysis approaches based on graph properties and topology. To solve this issue, we identified a list of "side compounds" corresponding to the main hub nodes of the network (S3 Table) that will be removed during the metabolic graph construction. To improve the performance of the graph theory methods, we also identified a list of reactions (S6 Table) to be excluded during the reaction metabolic graph construction. These reactions are blocked reactions (*i.e.*, reactions that cannot carry a flux in the GSMN), reactions always inactive in our cellular model, pool, sink, and exchange reactions. We used the metabolic reaction graph representation, where nodes are reactions and edges connect two reactions if the product of one is the substrate of the other. Metabolic reaction graphs were constructed according to the parameters defined previously before applying any graph-based methods such as the distance matrix calculation and subnetwork extractions.

Reaction distance matrix calculation

We computed the pairwise distance between reactions of interest by computing all the shortest paths between the nodes of the network corresponding to these reactions, resulting in a reaction pairwise distance matrix. To that purpose, we implemented a java app using the Met4j metabolic network toolbox (https://forgemia.inra.fr/metexplore/met4j). This app takes as input the GSMN (*i.e.*, Recon2.2*)* SBML file, a list of side compounds, a list of reactions to exclude, and a list of reactions between which the distances will be computed (*e.g.*, DARs in our case).

After loading the SBML file, the GSMN is pruned by removing side compounds and reactions to exclude. Then the JGraphT [64] ManytoMany shortest path implementation [65] is used. This implementation has been preferred over more common implementations such as Dijkstra [66] or Floyd–Warshall [67] due to its performance and its ability to take a specific list of reactions as input. The distance matrix is then saved in a comma separated file.

DARs clustering

For each DARs distance matrix computed with Met4J, we performed a hierarchical clustering with the Ward algorithm implemented in the SciPy python module. To control the size of clusters, we visualized each clustering result with a dendrogram representation and manually determined the number of clusters. Clusters were then obtained with the cutree function.

Subnetwork extraction

For each cluster of DARs, we extracted a minimal subnetwork with the minimal Steiner tree algorithm implemented in Met4J. This algorithm searches for a minimum spanning tree that contains the set of DARs and a minimum set of nodes from the GSMN to connect them according to the initial network topology.

Subnetwork visualization

We visualized the extracted subnetworks with MetExploreViz [39]. MetExploreViz is a JavaScript library dedicated to visualization for GSMNs that is integrated within the MetExplore [68] webserver. The list of metabolic reactions contained in the subnetwork is first mapped on the selected GSMN with MetExplore (https://metexplore.toulouse.inrae.fr/index.html/). From this mapping, we can then visualize and prune (*i.e.*, remove side compounds) the bipartite metabolic graph interactively with MetExploreViz. Two visualizations were done for each of the subnetworks. One with mapped up-activated DARs colored in red and down-activated DARs colored in green (Fig 4 A and 4 B), and a second one with reaction links colored according to the pathway associated with the reaction (Fig 4 B and 5 B). Interactive visualizations were saved online and are accessible via links provided in supplementary materials.

Workflow implementation

We partitioned the workflow into three jupyter notebooks with a common properties file. The first notebook, named "partial_enumeration.ipynb," takes as input barcode z-scores and generates batch files to execute the partial enumeration protocol on a SLURM computing cluster. Of note, these batch files can also be executed directly on a standard Linux (*i.e.*, Ubuntu, Debian, etc.) operating system. The second notebook, named "dars_calculation.ipynb," takes as input the partial-enumeration results from the previous notebook and computes baseline noise and DARs. Finally, the third notebook, named "analysis.ipynb," handles the network analysis step of our workflow, which includes computing the distance matrix for each list of DARs, hierarchical clustering, and extracting subnetworks with Met4J.

Python library code and jupyter notebooks are available on GitLab: https://forgemia.inra.fr/metexplore/MANA

Computing environment

Condition-specific modeling and partial enumeration requires CPLEX v20.10 to be installed on your system. Batch files generated by the "partial_enumeration.ipynb" notebook are designed to be launched on a SLURM computing grid but can also be launched on a standard Linux operating system. A local version of dexom-python is also required and can be cloned from the dexom-python repository (https://github.com/MetExplore/dexom-python). Finally, the network analysis workflow calling Met4J apps requires at least Java 11 to be installed on your machine.

Acknowledgements

We thank Dr. Hiroshi Yamada for providing supplementary data regarding the cell batches for the essays from the Open TG-GATEs database. We also thank Dr. Pablo Rodriguez-Mier for his expertise on how to tailor DEXOM to our needs and the discussions regarding condition-specific modeling and partial enumeration. Finally, we thank Dr. Gladys Ouedraogo for proofreading the manuscript and discussions throughout the project, and Dr. Bathilde Ambroise for the very interesting discussions we had during the very first phase of this work.

Funding

This work is supported by the Association Nationale de la Recherche et de la Technologie (ANRT) and by the French National Facility in Metabolomics and Fluxomics, MetaboHUB (11-INBS-0010), launched by the French Ministry of Research and Higher Education and the French ANR funding agency within the "Investissements d'Avenir" program.

Supplementary materials table of contents:

S1 Appendix Supplementary information. Pdf including: Correction coefficient formula, DAR specificity ratio formula, DEXOM parameters, Systematic sampling.

S1 Table DAR filtered out by the baseline noise filter. This table lists DAR having a R2 that is less than two times the noise estimate for this reaction, therefore filtered out by the baseline noise filter.

S2 Table DAR metrics. This table contains several metrics which aim at understanding how perturbed reactions are commonly perturbed between the eight studied chemicals. This table lists size, specificity ratio and average percentage of DARs intersection size.

S3 Table Side compounds identifiers. This table lists the sides compounds used in the network analysis stages.

S4 Table DAR subnetwork coverage. This table lists the DAR subnetwork coverage for each identified subnetwork for amiodarone and valproic acid.

S5 Table DEG lists size obtained from the ToxicoDB database. This table summarize key metadata for each compound such as the dose, exposure time, number of replicates, cell type, cell batch, control vehicle, DEG signature size and metabolic DEG signature size. Were considered differentially expressed, genes having a log2(abs(FC)) above 0.26 and a FDR corrected p-value below 0.05.

S6 Table Excluded reactions identifiers. This table lists the reactions that have been excluded in the network analysis stages.

S1 Fig Pathway enrichment performed on DARs with MetExploreViz. Pathway over-representation analysis performed with a Fisher Exact test; p-values corrected with the Benjamini-Hochberg method on Recon2.2 metabolic pathways for 8 hepatotoxic molecules.

S2 Fig Pathway enrichment performed on valproic acid DEGs with ReactomePA. Pathway over-representation analysis performed with a Fisher Exact test, p-values corrected with the Benjamini-Hochberg method on Reactome 2022 pathways (genes with log2(absFC)) > 0.26 and FDR-corrected p-values < 0.05 were considered as DEG)

References

1. Alexander-White C, Bury D, Cronin M, Dent M, Hack E, Hewitt NJ, et al. A 10-step framework for use of read-across (RAX) in next generation risk assessment (NGRA) for cosmetics safety assessment. Regul Toxicol Pharmacol [Internet]. 2022;129:105094. Available from: https://www.sciencedirect.com/science/article/pii/S027323002100235X
2. Li T, Tong W, Roberts R, Liu Z, Thakkar S. Deep Learning on High-Throughput Transcriptomics to Predict Drug-Induced Liver Injury [Internet]. Vol. 8, Frontiers in Bioengineering and Biotechnology . 2020. p. 1366. Available from: https://www.frontiersin.org/article/10.3389/fbioe.2020.562677
3. Harrill JA, Everett LJ, Haggard DE, Sheffield T, Bundy JL, Willis CM, et al. High-Throughput Transcriptomics Platform for Screening Environmental Chemicals. Toxicol Sci. 2021 Apr;181(1):68–89.
4. Chen Q, Wu L, Liu W, Xing L, Fan X. Enhanced QSAR Model Performance by Integrating Structural and Gene Expression Information. Vol. 18, Molecules. 2013. p. 10789–801.
5. Heindel JJ, Blumberg B, Cave M, Machtinger R, Mantovani A, Mendez MA, et al. Metabolism disrupting chemicals and metabolic disorders. Reprod Toxicol. 2017 Mar;68:3–33.
6. Sarni ROS, Kochi C, Suano-Souza FI. Childhood obesity: an ecological perspective. J Pediatr (Rio J). 2022;98 Suppl 1(Suppl 1):S38–46.
7. Gore AC, Chappell VA, Fenton SE, Flaws JA, Nadal A, Prins GS, et al. EDC-2: The Endocrine Society's Second Scientific Statement on Endocrine-Disrupting Chemicals. Endocr Rev. 2015 Dec;36(6):E1–150.
8. Heal DJ, Gosden J, Jackson HC, Cheetham SC, Smith SL. Metabolic consequences of antipsychotic therapy: preclinical and clinical perspectives on diabetes, diabetic ketoacidosis, and obesity. Handb Exp Pharmacol. 2012;(212):135–64.
9. Miranda RA, Silva BS, de Moura EG, Lisboa PC. Pesticides as endocrine disruptors: programming for obesity and diabetes. Endocrine. 2023 Mar;79(3):437–47.

10. Nguyen T-M, Shafi A, Nguyen T, Draghici S. Identifying significantly impacted pathways: a comprehensive review and assessment. Genome Biol [Internet]. 2019;20(1):203. Available from: https://doi.org/10.1186/s13059-019-1790-4

11. Wieder C, Frainay C, Poupin N, Rodríguez-Mier P, Vinson F, Cooke J, et al. Pathway analysis in metabolomics: Recommendations for the use of over-representation analysis. PLoS Comput Biol. 2021 Sep;17(9):e1009105.

12. Karp PD, Midford PE, Caspi R, Khodursky A. Pathway size matters: the influence of pathway granularity on over-representation (enrichment analysis) statistics. BMC Genomics [Internet]. 2021;22(1):191. Available from: https://doi.org/10.1186/s12864-021-07502-8

13. Haider S, Black MB, Parks BB, Foley B, Wetmore BA, Andersen ME, et al. A Qualitative Modeling Approach for Whole Genome Prediction Using High-Throughput Toxicogenomics Data and Pathway-Based Validation [Internet]. Vol. 9, Frontiers in Pharmacology . 2018. Available from: https://www.frontiersin.org/articles/10.3389/fphar.2018.01072

14. Liu Z, Zhu L, Thakkar S, Roberts R, Tong W. Can Transcriptomic Profiles from Cancer Cell Lines Be Used for Toxicity Assessment? Chem Res Toxicol [Internet]. 2020 Jan 21;33(1):271–80. Available from: https://doi.org/10.1021/acs.chemrestox.9b00288

15. Soufan O, Ewald J, Viau C, Crump D, Hecker M, Basu N, et al. T1000: a reduced gene set prioritized for toxicogenomic studies. PeerJ. 2019;7:e7975.

16. Mav D, Shah RR, Howard BE, Auerbach SS, Bushel PR, Collins JB, et al. A hybrid gene selection approach to create the S1500+ targeted gene sets for use in high-throughput transcriptomics. PLoS One [Internet]. 2018 Feb 20;13(2):e0191105. Available from: https://doi.org/10.1371/journal.pone.0191105

17. Ganter B, Snyder RD, Halbert DN, Lee MD. Toxicogenomics in drug discovery and development: mechanistic analysis of compound/class-dependent effects using the DrugMatrix database. Pharmacogenomics. 2006 Oct;7(7):1025–44.

18. Lamb J, Crawford ED, Peck D, Modell JW, Blat IC, Wrobel MJ, et al. The connectivity map: Using gene-expression signatures to connect small molecules, genes, and disease. Science (80-) [Internet]. 2006 Sep 29 [cited 2020 Feb 13];313(5795):1929–35. Available from: https://www.sciencemag.org/lookup/doi/10.1126/science.1132939

19. Subramanian A, Narayan R, Corsello SM, Peck DD, Natoli TE, Lu X, et al. A Next Generation Connectivity Map: L1000 Platform and the First 1,000,000 Profiles. Cell. 2017 Nov;171(6):1437-1452.e17.

20. Nair SK, Eeles C, Ho C, Beri G, Yoo E, Tkachuk D, et al. ToxicoDB: an integrated database to mine and visualize large-scale toxicogenomic datasets. Nucleic Acids Res [Internet]. 2020 May 18;48(W1):W455–62. Available from: https://doi.org/10.1093/nar/gkaa390

21. Igarashi Y, Nakatsu N, Yamashita T, Ono A, Ohno Y, Urushidani T, et al. Open TG-GATEs: a large-scale toxicogenomics database. Nucleic Acids Res [Internet]. 2014/10/13. 2015 Jan;43(Database issue):D921–7. Available from: https://pubmed.ncbi.nlm.nih.gov/25313160

22. Heusinkveld HJ, Wackers PFK, Schoonen WG, van der Ven L, Pennings JLA, Luijten M. Application of the comparison approach to open TG-GATEs: A useful toxicogenomics tool for detecting modes of action in chemical risk assessment. Food Chem Toxicol an Int J Publ Br Ind Biol Res Assoc. 2018 Nov;121:115–23.

23. Terzer M, Maynard ND, Covert MW, Stelling J. Genome-scale metabolic networks. WIREs Syst Biol Med [Internet]. 2009 Nov 1;1(3):285–97. Available from: https://doi.org/10.1002/wsbm.37

24. Thiele I, Palsson BØ. A protocol for generating a high-quality genome-scale metabolic reconstruction. Nat Protoc [Internet]. 2010/01/07. 2010 Jan;5(1):93–121. Available from: https://pubmed.ncbi.nlm.nih.gov/20057383

25. Thiele I, Swainston N, Fleming RMT, Hoppe A, Sahoo S, Aurich MK, et al. A community-driven global reconstruction of human metabolism. Nat Biotechnol [Internet]. 2013/03/03. 2013 May;31(5):419–25. Available from: https://pubmed.ncbi.nlm.nih.gov/23455439

26. Swainston N, Smallbone K, Hefzi H, Dobson PD, Brewer J, Hanscho M, et al. Recon 2.2: from reconstruction to model of human metabolism. Metabolomics. 2016;12:109.

27. Brunk E, Sahoo S, Zielinski DC, Altunkaya A, Dräger A, Mih N, et al. Recon3D enables a three-dimensional view of gene variation in human metabolism. Nat Biotechnol. 2018 Mar;36(3):272–81.

28. Robinson JL, Kocabaş P, Wang H, Cholley P-E, Cook D, Nilsson A, et al. An atlas of human metabolism. Sci Signal. 2020 Mar;13(624).

29. Zur H, Ruppin E, Shlomi T. iMAT: an integrative metabolic analysis tool. Bioinformatics. 2010 Dec;26(24):3140–2.

30. Shlomi T, Cabili MN, Herrgård MJ, Palsson BØ, Ruppin E. Network-based prediction of human tissue-specific metabolism. Nat Biotechnol. 2008 Sep;26(9):1003–10.

31. Pusa T, Ferrarini MG, Andrade R, Mary A, Marchetti-Spaccamela A, Stougie L, et al. MOOMIN – Mathematical explOration of 'Omics data on a MetabolIc Network. Bioinformatics [Internet]. 2020 Jan 15;36(2):514–23. Available from: https://doi.org/10.1093/bioinformatics/btz584

32. Vlassis N, Pacheco MP, Sauter T. Fast Reconstruction of Compact Context-Specific Metabolic Network Models. PLOS Comput Biol [Internet]. 2014 Jan 16;10(1):e1003424. Available from: https://doi.org/10.1371/journal.pcbi.1003424

33. Agren R, Bordel S, Mardinoglu A, Pornputtapong N, Nookaew I, Nielsen J. Reconstruction of Genome-Scale Active Metabolic Networks for 69 Human Cell Types and 16 Cancer Types Using INIT. PLOS Comput Biol [Internet]. 2012 May 17;8(5):e1002518. Available from: https://doi.org/10.1371/journal.pcbi.1002518

34. Rodríguez-Mier P, Poupin N, de Blasio C, Le Cam L, Jourdan F. DEXOM: Diversity-based enumeration of optimal context-specific metabolic networks. PLOS Comput Biol [Internet]. 2021 Feb 11;17(2):e1008730. Available from: https://doi.org/10.1371/journal.pcbi.1008730

35. Rossell S, Huynen MA, Notebaart RA. Inferring Metabolic States in Uncharacterized Environments Using Gene-Expression Measurements. PLOS Comput Biol [Internet]. 2013 Mar 21;9(3):e1002988. Available from: https://doi.org/10.1371/journal.pcbi.1002988

36. Robaina-Estévez S, Nikoloski Z. On the effects of alternative optima in context-specific metabolic model predictions. PLOS Comput Biol [Internet]. 2017 May 30;13(5):e1005568. Available from: https://doi.org/10.1371/journal.pcbi.1005568

37. Urushidani T. Prediction of Hepatotoxicity Based on the Toxicogenomics Database [Internet]. Hepatotoxicity. 2007. p. 507–29. (Wiley Online Books). Available from: https://doi.org/10.1002/9780470516751.ch20

38. Noriyuki N, Igarashi Y, Ono A, Yamada H, Ohno Y, Urushidani T. Evaluation of DNA microarray results in the Toxicogenomics Project (TGP) consortium in Japan. J Toxicol Sci. 2012;37(4):791–801.

39. Chazalviel M, Frainay C, Poupin N, Vinson F, Merlet B, Gloaguen Y, et al. MetExploreViz: web component for interactive metabolic network visualization. Bioinformatics [Internet]. 2018 Jan 15;34(2):312–3. Available from: https://pubmed.ncbi.nlm.nih.gov/28968733

40. Anthérieu S, Rogue A, Fromenty B, Guillouzo A, Robin M-A. Induction of vesicular steatosis by amiodarone and tetracycline is associated with up-regulation of lipogenic genes in heparg cells. Hepatology [Internet]. 2011 Jun 1;53(6):1895–905. Available from: https://doi.org/10.1002/hep.24290

41. Mnif L, Sellami R, Masmoudi J. Valproic Acid and Hepatic Steatosis: A Possible Link? About a Case Report. Psychopharmacol Bull. 2016 Aug;46(2):59–62.

42. Poupin N, Corlu A, Cabaton NJ, Dubois-Pot-Schneider H, Canlet C, Person E, et al. Large-Scale Modeling Approach Reveals Functional Metabolic Shifts during Hepatic Differentiation. J Proteome Res [Internet]. 2019 Jan 4;18(1):204–16. Available from: https://doi.org/10.1021/acs.jproteome.8b00524

43. Allard J, Bucher S, Massart J, Ferron P-J, Le Guillou D, Loyant R, et al. Drug-induced hepatic steatosis in absence of severe mitochondrial dysfunction in HepaRG cells: proof of multiple mechanism-based toxicity. Cell Biol Toxicol [Internet]. 2021;37(2):151–75. Available from: https://doi.org/10.1007/s10565-020-09537-1

44. Lee C, Cunningham P. Community detection: effective evaluation on large social networks. J Complex Networks [Internet]. 2014 Mar 1;2(1):19–37. Available from: https://doi.org/10.1093/comnet/cnt012

45. Karp RM. Reducibility among Combinatorial Problems BT - Complexity of Computer Computations: Proceedings of a symposium on the Complexity of Computer Computations, held March 20–22, 1972, at the IBM Thomas J. Watson Research Center, Yorktown Heights, New York, an. In: Miller RE, Thatcher JW, Bohlinger JD, editors. Boston, MA: Springer US; 1972. p. 85–103. Available from: https://doi.org/10.1007/978-1-4684-2001-2_9

46. Voß S. Steiner's problem in graphs: heuristic methods. Discret Appl Math [Internet]. 1992;40(1):45–72. Available from: https://www.sciencedirect.com/science/article/pii/0166218X92900212

47. Göttlicher M, Minucci S, Zhu P, Krämer OH, Schimpf A, Giavara S, et al. Valproic acid defines a novel class of HDAC inhibitors inducing differentiation of transformed cells. EMBO J. 2001 Dec;20(24):6969–78.

48. Zhou J, Wang X, Wang M, Chang Y, Zhang F, Ban Z, et al. The lysine catabolite saccharopine impairs development by disrupting mitochondrial homeostasis. J Cell Biol [Internet]. 2018 Dec 20;218(2):580–97. Available from: https://doi.org/10.1083/jcb.201807204

49. Fromenty B, Pessayre D. Inhibition of mitochondrial beta-oxidation as a mechanism of hepatotoxicity. Pharmacol Ther. 1995;67(1):101–54.

50. Lheureux PER, Penaloza A, Zahir S, Gris M. Science review: Carnitine in the treatment of valproic acid-induced

toxicity – what is the evidence? Crit Care [Internet]. 2005;9(5):431. Available from: https://doi.org/10.1186/cc3742

51. Maciejak P, Szyndler J, Kołosowska K, Turzyńska D, Sobolewska A, Walkowiak J, et al. Valproate disturbs the balance between branched and aromatic amino acids in rats. Neurotox Res. 2014 May;25(4):358–68.

52. Shibata K, Kondo R, Sano M, Fukuwatari T. Increased conversion of tryptophan to nicotinamide in rats by dietary valproate. Biosci Biotechnol Biochem. 2013;77(2):295–300.

53. Hubel E, Fishman S, Holopainen M, Käkelä R, Shaffer O, Houri I, et al. Repetitive amiodarone administration causes liver damage via adipose tissue ER stress-dependent lipolysis, leading to hepatotoxic free fatty acid accumulation. Am J Physiol Liver Physiol [Internet]. 2021 Jul 14;321(3):G298–307. Available from: https://doi.org/10.1152/ajpgi.00458.2020

54. Gautier L, Cope L, Bolstad BM, Irizarry RA. affy—analysis of Affymetrix GeneChip data at the probe level. Bioinformatics [Internet]. 2004 Feb 12;20(3):307–15. Available from: https://doi.org/10.1093/bioinformatics/btg405

55. Irizarry RA, Hobbs B, Collin F, Beazer-Barclay YD, Antonellis KJ, Scherf U, et al. Exploration, normalization, and summaries of high density oligonucleotide array probe level data. Biostatistics. 2003 Apr;4(2):249–64.

56. Leek JT, Johnson WE, Parker HS, Jaffe AE, Storey JD. The sva package for removing batch effects and other unwanted variation in high-throughput experiments. Bioinformatics [Internet]. 2012/01/17. 2012 Mar 15;28(6):882–3. Available from: https://pubmed.ncbi.nlm.nih.gov/22257669

57. Carlson M. hgu133plus2.db: Affymetrix Human Genome U133 Plus 2.0 Array annotation data. 2016.

58. Pagès H, Carlson M, Falcon S, Li N. AnnotationDbi: Manipulation of SQLite-based annotations in Bioconductor. 2020.

59. McCall MN, Bolstad BM, Irizarry RA. Frozen robust multiarray analysis (fRMA). Biostatistics. 2010 Apr;11(2):242–53.

60. McCall MN, Uppal K, Jaffee HA, Zilliox MJ, Irizarry RA. The Gene Expression Barcode: leveraging public data repositories to begin cataloging the human and murine transcriptomes. Nucleic Acids Res. 2011 Jan;39(Database issue):D1011-5.

61. McCall MN, Jaffee HA, Zelisko SJ, Sinha N, Hooiveld G, Irizarry RA, et al. The Gene Expression Barcode 3.0: improved data processing and mining tools. Nucleic Acids Res. 2014 Jan;42(Database issue):D938-43.

62. Guguen-Guillouzo C, Guillouzo A. General Review on In Vitro Hepatocyte Models and Their Applications BT - Hepatocytes: Methods and Protocols. In: Maurel P, editor. Totowa, NJ: Humana Press; 2010. p. 1–40. Available from: https://doi.org/10.1007/978-1-60761-688-7_1

63. Frainay C, Jourdan F. Computational methods to identify metabolic sub-networks based on metabolomic profiles. Brief Bioinform. 2017 Jan;18(1):43–56.

64. Michail D, Kinable J, Naveh B, Sichi J V. JGraphT—A Java Library for Graph Data Structures and Algorithms. ACM Trans Math Softw [Internet]. 2020;46(2). Available from: https://doi.org/10.1145/3381449

65. Knopp S, Sanders P, Schultes D, Schulz F, Wagner D. Computing Many-to-Many Shortest Paths Using Highway Hierarchies. In: 2007 Proceedings of the Workshop on Algorithm Engineering and Experiments (ALENEX) [Internet]. Society for Industrial and Applied Mathematics; 2007. p. 36–45. (Proceedings). Available from: https://doi.org/10.1137/1.9781611972870.4

66. Dijkstra EW. A note on two problems in connexion with graphs. Numer Math [Internet]. 1959;1(1):269–71. Available from: https://doi.org/10.1007/BF01386390

67. Floyd RW. Algorithm 97: Shortest Path. Commun ACM [Internet]. 1962;5(6):345. Available from: https://doi.org/10.1145/367766.368168

68. Cottret L, Frainay C, Chazalviel M, Cabanettes F, Gloaguen Y, Camenen E, et al. MetExplore: collaborative edition and exploration of metabolic networks. Nucleic Acids Res [Internet]. 2018 Jul 2;46(W1):W495–502. Available from: https://doi.org/10.1093/nar/gky301