

Gene-metabolite annotation with shortest reactional distance enhances metabolite genome-wide association studies results

Cantin Baron ^{1,2}, Sarah Cherkaoui ^{2,3,4}, Sandra Therrien-Laperriere ², Yann Ilboudo ^{1,2}, Raphaël Poujol ², Pamela Mehanna ², Melanie E. Garrett ⁵, Marilyn J. Telen⁶, Allison E. Ashley-Koch ⁵, Pablo Bartolucci^{7,8}, John D. Rioux ^{1,2,9}, Guillaume Lettre ^{2,9}, Christine Des Rosiers ^{1,2,10}, Matthieu Ruiz ^{2,10}, Julie G. Hussin^{2,9, *}

¹ Département de Biochimie et de Médecine Moléculaire, Université de Montréal, Québec, Canada

² Montreal Heart Institute, Québec, Canada

³ Division of Oncology and Children's Research Center, University Children's Hospital Zurich, University of Zurich, Switzerland

⁴ Department of Pediatric and Adolescent Oncology, Gustave Roussy Cancer Center, Université Paris-Saclay, Villejuif, France

⁵ Duke Molecular Physiology Institute, Duke University Medical Center, Durham, NC, USA

⁶ Division of Hematology, Department of Medicine, Duke University Medical Center, Durham, NC, USA

⁷ Université Paris Est Créteil, Hôpitaux Universitaires Henri Mondor, APHP, Sickle cell referral center – UMGGR, Créteil, France

⁸ Université Paris Est Créteil, IMRB, Laboratory of excellence LABEX, Créteil, France

⁹ Département de Médecine, Université de Montréal, Québec, Canada

¹⁰ Département de Nutrition, Université de Montréal, Québec, Canada

* Correspondence: julie.hussin@umontreal.ca

SUMMARY

1 Studies combining metabolomics and genetics, known as metabolite genome-wide association
2 studies (mGWAS), have provided valuable insights into our understanding of the genetic control
3 of metabolite levels. However, the biological interpretation of these associations remains
4 challenging due to a lack of existing tools to annotate mGWAS gene-metabolite pairs beyond the
5 use of conservative statistical significance threshold. Here, we computed the shortest reactional
6 distance (SRD) based on the curated knowledge of the KEGG database to explore its utility in
7 enhancing the biological interpretation of results from three independent mGWAS, including a
8 case study on sickle cell disease patients. Results show that, in reported mGWAS pairs, there is an
9 excess of small SRD values and that SRD values and p-values significantly correlate, even beyond
10 the standard conservative thresholds. The added-value of SRD annotation is shown for
11 identification of potential false negative hits, exemplified by the finding of gene-metabolite
12 associations with $SRD \leq 1$ that did not reach standard genome-wide significance cut-off. The wider
13 use of this statistic as an mGWAS annotation would prevent the exclusion of biologically relevant
14 associations and can also identify errors or gaps in current metabolic pathway databases. Our
15 findings highlight the SRD metric as an objective, quantitative and easy-to-compute annotation
16 for gene-metabolite pairs that can be used to integrate statistical evidence to biological networks.
17

18 **Keywords:** mGWAS, KEGG, biological networks, shortest reactional distance, annotation

19

20

21 INTRODUCTION

22 High throughput biotechnologies and analytic approaches applied to large human cohorts have
23 recently revolutionized biomedical research, allowing the quantification and characterization of
24 biological molecules to generate “omics” datasets. Genomics, the characterization of an
25 individual’s DNA molecules, led to the identification of thousands of genetic variants associated
26 with a trait, disease, or response to treatments, through large-scale genome-wide association
27 studies (GWAS) [1]. Although GWAS have resulted in a better understanding of disease
28 mechanisms, these approaches only consider genetic variation established at birth, and ignore the
29 environment of an individual, influencing its biological state. An alternative strategy to
30 complement traditional GWAS and better understand human biology is to comprehensively
31 interrogate disease states at the molecular level using metabolomics, which offers a robust way to
32 systematically measure thousands of low-molecular-weight compounds, called metabolites. After
33 tissue extraction or collection of biological samples (usually blood and urine), metabolites can be
34 detected, identified, and quantified using either mass spectrometry (MS) or nuclear magnetic
35 resonance (NMR) [2]. As biomarkers of the underlying molecular dysfunctions, metabolite levels
36 correspond to intermediate phenotypes (or endophenotypes) representing natural or clinical
37 heterogeneity. Storing and sharing metabolomics knowledge is the focus of The Human
38 Metabolome Database (HMDB), which is one of the largest and comprehensive curated collection
39 of human metabolite and human metabolism data in the world [3].

40 Metabolomics can be seen as the study of the ultimate molecular response of an organism to
41 genetic, environmental, and pathological modifications, but elucidating specific molecular
42 mechanisms from metabolomics alone is difficult since many metabolites are involved in multiple
43 biological processes. Linking metabolomics signals with genetics provides a promising approach
44 to identify the implicated pathways and, as such, metabolite genome-wide association studies
45 (mGWAS) are key approaches to integrate metabolomics with genomics. mGWAS report
46 associations between metabolites and genetic loci, referred to as metabolite quantitative trait loci
47 (mQTL), making it possible to study associations between millions of genetic variants and
48 thousands of metabolites and to generate insightful hypotheses about uncharacterized regulatory
49 mechanisms [2]. Although very powerful, interpretation of mGWAS results remains challenging.
50 Indeed, millions of statistical tests are generally performed between single nucleotide
51 polymorphisms (SNPs) and metabolites, leading to a huge multiple testing burden [4].

52 Furthermore, interpreting the biological meaning of an association between a given SNP and a
53 metabolite requires interdisciplinary knowledge. In this context, the simplest way of prioritizing
54 hypotheses remains to apply a conservative correction for multiple testing such as Bonferroni or
55 Benjamini-Hochberg procedures [4]. The consequence of this correction is that only associations
56 with the most significant p-values are reported whereas biologically relevant associations with a
57 suggestive p-value may be missed.

58 With the increasing number of published mGWAS and their current limits, it is necessary to
59 develop systematic methodologies to gain better insight into the mGWAS results by exploiting
60 the most up-to-date biological knowledge. There are multiple resources available that aim at
61 storing and describing known relationships between genes and phenotypes or endophenotypes
62 (GWAS Catalog [5], PhenoScanner [6], OpenGWAS [7], Open Targets Genetics [8], PheLiGe
63 [9], DisGeNET [10]) as well as tools to annotate gene-metabolites pairs based on the available
64 resources but none of them have been specifically designed to address current mGWAS
65 limitations [11, 12]. Only few methods have been developed to gain insight into the biological
66 interpretation of mGWAS data specifically [13, 14], and none have been explicitly used to
67 annotate mGWAS results based on well-curated biochemical knowledge, such as KEGG database
68 [15]. KEGG is recognized for its high curation level of metabolic pathways for many model
69 organisms, including humans, and is a reference for the reporting of enzymatic reactions, which
70 is needed to understand the functional importance of gene-metabolite pairs. KEGG is a reference
71 for the biochemical functions of genes with the descriptions of enzyme reactions, but a systematic
72 and quantitative annotation procedure using this database in the context of mGWAS results
73 remains to be explored. In this regard, recent studies have highlighted the applicability and utility
74 of topological data analysis based on graph theory approaches [16], including the shortest path
75 [17, 18], in revealing biological knowledge in the context of complex metabolic networks [13,
76 17, 18]. For example, the shortest path was used to characterize the impact of gene deletion on
77 nearby metabolites within the metabolic network of *E. coli* [18]. It was also used to analyze the
78 relationship between expression quantitative trait loci (eQTL) and metabolomic data, for example
79 within the metabolic network of rat adipose tissues [17]. Still, this type of metric has not been
80 applied to systematically annotate mGWAS results.

81 In this study, we assessed the utility of the shortest reactional distance (SRD) metric computed
82 from KEGG database's pathways to annotate mGWAS results and to help extract biological

83 insights from them. We developed PathQuant, an R package, to enable a robust and systematic
84 computation of SRD values between any lists of gene-metabolite pairs mapped onto KEGG
85 graphs, which represent metabolic pathways, while keeping the original, well-curated, topology
86 of each queried pathway. Focusing on genes encoding for enzymes and their associated
87 metabolites, we applied the SRD annotation to two previously published mGWAS datasets in
88 individuals from different ethnicities: an mGWAS [11] performed on 7,824 participants from the
89 TwinsUK cohort [19] and the KORA study [20], referred herein as the TK study, and an mGWAS
90 [21] performed on 614 Qatari participants from the Hamad Medical Corporation (HMC), referred
91 herein as the HMC study. We explored results at varying levels of statistical significance. We
92 found that the SRD metric enables identification of associations that do not meet currently
93 accepted cut-off of statistical significance (suggestive associations) but which have high
94 biological relevance. Finally, we performed an mGWAS on previously reported genetic and
95 metabolomic data from a Sickle Cell Disease (SCD) cohort [22, 23], referred herein as the SCD
96 study, and show how the SRD metric can be used to prioritize novel hits, including hits with a
97 suggestive p-value.

98

99 **RESULTS**

100 **Overview of study pipeline**

101 PathQuant is a tool that converts a metabolic pathway map into a graph of biochemical reactions
102 with metabolites as nodes and genes as edges, to compute the SRD path between a given gene-
103 metabolite pair (STAR Methods, Supplementary Methods). To explore the potential of SRD
104 values as an annotation metric to inform on the biological relevance of gene-metabolite pairs
105 obtained from mGWAS, we developed a pipeline, presented in Figure 1. First, we gather
106 mGWAS summary statistics and perform standard quality-control filters on genetic variants
107 based on minor allele frequencies (MAF) and completeness. The second step of our pipeline only
108 keeps bi-allelic SNPs at $MAF > 0.01$ (to exclude rare variants) and tested across all measured
109 metabolites, but indels, tri-allelic SNPs and rare variants can be easily integrated, if present in
110 mGWAS summary statistics. Third, SNPs are mapped to their closest genes coding for an enzyme
111 (within 10,000 pb upstream and downstream) to obtain gene-metabolite pairs. As multiple SNPs
112 will generally be linked with the same gene, only the minimum p-value of all SNPs is kept for a
113 gene-metabolite pair, representing the strongest statistical signal for each pair within a study.
114 These p-values can be categorized as either significant, suggestive, or non-significant according
115 to appropriate thresholds that depend on the dataset (STAR Methods). We then retrieve KEGG
116 IDs from all pairs (STAR Methods), and only retain pairs for which both IDs could be found.
117 Third, we run PathQuant R package to compute the SRD value on all remaining pairs based on
118 the KEGG overview graph (hsa01100). Note that other KEGG graph, or a combination of KEGG
119 graphs, can be easily added at this step. SRD values obtained can be numerical when a path is
120 found, can be given an infinite value (noted as “Inf”) when there is no path between the gene and
121 metabolite in the queried graph, or an “NA” value when the gene or the metabolite (or both) are
122 not found in the queried graph, despite having a valid KEGG ID. We can then perform several
123 analyses of SRD values, as detailed below.

124

125 **Stringent and suggestive associated pairs have shorter reactional distances**

126 To test whether the SRD is a metric capable of capturing the biological relevance of a gene-
127 metabolite pairs, we first explored the results of the TK study, as it represents one of the largest
128 mGWAS studies conducted to date. This mGWAS was done on a group of participants from

129 European descent, and many findings were replicated in follow-up studies, making it a well-
130 validated dataset. Furthermore, it is considered a reference study by the community, resulting in
131 an ideal dataset to test the hypothesis that stringently associated pairs will have lower SRD values.
132 We started to explore this dataset by using the 74 gene-metabolite pairs with KEGG IDs passing
133 the genome-wide significance cut-off set by the original authors [11] (Table 1). Among these, 40
134 gene-metabolite pairs were mapped onto KEGG overview graph. After excluding the two pairs
135 with infinite SRD values, the median SRD value for the remaining 38 pairs is 1, which indicates
136 a close biological relationship between the genes and their associated metabolites (Figure 2). To
137 assess how significant this result is, we used two strategies: we compared the SRD values to (1)
138 a null distribution (STAR Methods), built from all possible pairs of gene-metabolites that are
139 present on the KEGG overview graph (Figure 2A); (2) to a permuted set of gene-metabolite pairs
140 (Figure 2B), built with the 33 metabolites and 27 genes involved in the 40 significant gene-
141 metabolite pairs (Figure 2C). There is a statistically significant difference between the TK pairs'
142 SRD values and the null distribution (Welch test, p-value < 3.48×10^{-16}), confirming the close
143 biological relationship between mGWAS gene-metabolite pairs in the TK study. Similarly, the
144 median SRD value for the permuted pairs is 8, which is significantly higher than the median SRD
145 of 1 (empirical p-value = 0.035). These results confirm that closely connected genes and
146 metabolites are enriched in gene-metabolite pairs discovered in mGWAS.

147 To visualize the relationship between all genes and metabolites involved in the significant pairs
148 reported in the TK study within KEGG overview graph, we used a heatmap representation of
149 SRD values (Figure 2C), highlighting the reported significant pairs using thick black boxes.
150 Seventeen of these associations have an SRD of 0, implying that these gene-metabolite pairs are
151 from a single enzymatic reaction. For example, PSPH (phosphoserine phosphatase) is catalysing
152 the formation of L-Serine and results in an SRD of 0. Interestingly, within the 40 associations
153 with an SRD annotation, some genes are closely connected to multiple metabolites, such as
154 ACADM and CPS1 that have multiple small SRD values. Moreover, some genes are isolated and
155 appear to be part of disconnected subgraphs of the KEGG overview graph, such as GMPR
156 (Guanosine Monophosphate Reductase), NAT2 (N-Acetyltransferase 2) and TDO2 (Tryptophan
157 2,3-Dioxygenase), which obtained infinite SRD values with most metabolites.

158 The design of mGWAS can be heterogeneous, with distinct genomic and metabolomic
159 technologies and pre-processing steps used across studies. To demonstrate the applicability of

160 our annotation more broadly, it is important to use an independent study to replicate the
161 observations above. Furthermore, there is a widely recognized bias towards white Europeans in
162 genetics studies, which often makes results less generalizable in other ethnicities [24, 25],
163 highlighting a need for new bioinformatics solutions to be tested in underrepresented populations.
164 In line with these criteria, we further tested our approach in the HMC study, which has been
165 performed in a different ethnic group, on Qatari individuals. In this mGWAS, the genomic data
166 comes from the exome sequencing technology compared to genotyping data in the TK study, and
167 the metabolomics data is generated using differing pre-processing steps (see STAR Methods).
168 Furthermore, the authors of this study identified SNP-metabolite pairs with suggestive p-values,
169 which were made available (Supplementary File number 7 of [21]), allowing us to explore
170 whether including the suggestive mGWAS association results replicated the observation from the
171 TK study. Of 68 gene-metabolite associations (Table 1) at a cut-off of $p \leq 1.4 \times 10^{-7}$, 42 were
172 mapped to the KEGG overview graph for which SRD values were computed. Four pairs had
173 infinite SRD values (see Supplementary Figure 1C), while the 38 gene-metabolite remaining
174 pairs (Table 1) had SRDs of 0 or 1 (Supplementary Figure 1), meaning that we have associations
175 that are either substrate-product associations or with one intermediate step. Thus, we replicated
176 the observation seen in the TK study of a significant enrichment of low SRD values compared to
177 the null distribution in this second mGWAS (Welch test, p-value $< 7.91 \times 10^{-60}$, Supplementary
178 Figure 1), even when considering a lower significant threshold than standard genome-wide cut-
179 offs.

180

181 **SRD annotation can identify false negative hits**

182 By using two different studies and different summary statistics cut-offs, we have determined that
183 associations with stringent and suggestive p-values are enriched for low SRDs. We next explored
184 more formally the relationship between the SRD values and the p-values from the TK study. We
185 graphically defined the gene-metabolite significance cut-off at $p < 3.16 \times 10^{-4}$ of the TK study
186 based on a QQplot (Supplementary Figure 2B, STAR Methods) and used it as a threshold for the
187 minimum value to determine the relationship between SRDs and p-values. We observed a
188 significant negative correlation between the mGWAS p-values and SRD values ($R = -0.2$, $p =$
189 9.1×10^{-11} , Figure 3A), meaning that the higher the significance of an association between a

190 metabolite and a gene is, the closer they are likely to be in the KEGG overview graph. This
191 negative correlation confirms the enrichment of biologically relevant pairs in significant
192 associations, and suggests that combining p-values and SRD may help prioritizing associations
193 of high biological interest. Furthermore, the SRD value could be used to select for further analysis
194 pairs with small SRD having low p-value that nevertheless do not pass stringent significance cut-
195 offs.

196 As a result of multiple testing burden between each variant and each metabolite level, only top
197 hits are generally reported in mGWAS, but it is well known that there could be false negative hits
198 [26]. Thereby, we tested the potential for the SRD metric to identify some of these false-negative
199 candidates in these published mGWAS. In the TK study, we focused on the genes and metabolites
200 involved in stringent associations (presented in the heatmap, Figure 2C), which include 27 genes
201 and 33 metabolites. We noticed gene-metabolite pairs annotated with a smaller SRD than the
202 initially reported association, such as the ALDH18A1 (aldehyde dehydrogenase 18 family
203 member A1) and L-citrulline pair, which has an SRD value of 2, whereas other pairs involving
204 L-citrulline show smaller SRD values. For instance, the CPS1 (carbamoyl-phosphate synthase 1)
205 and L-citrulline pair has an SRD = 1 but did not pass the stringent p-value cut-offs of the TK
206 study, despite a p-value of 1.749×10^{-8} . The SRD value of 1 reflects a sequence of reactions
207 pertaining to the urea cycle catalysed by CPS1 for the formation carbamoyl phosphate from
208 ammonia and bicarbonate and by the ornithine transcarbamylase (OTC) for the formation of L-
209 citrulline from carbamoyl phosphate and L-ornithine [27]. The association between CPS1 and L-
210 citrulline has recently been described twice in the literature with a p-value of 1×10^{-25} with SNP
211 rs1509820 [28] and 1×10^{-14} with SNP rs975530777 [29]. The fact that this association was not
212 significant in the TK study can thus be considered a false negative result.

213 Using the same approach, we also identified an example within the HMC study of a gene-
214 metabolite pair at SRD = 0, which was not reported as a significant association because of a
215 suggestive p-value ($p = 1.17 \times 10^{-7}$). This association is between UMPS (uridine monophosphate
216 synthetase) and orotate (Supplementary Figure 1C, indicated in red). UMPS is a bifunctional
217 enzyme that is part of the *de novo* pyrimidine biosynthetic pathway; its orotate
218 phosphoribosyltransferase subunit catalyses the addition of ribose-5-phosphate to orotate to form
219 orotidine monophosphate (KEGG reaction ID R01870). Given that a significant association has

220 been previously reported between UMPS gene and orotate [30], we could consider that this
221 finding is replicated in the HMC study but only by adding information about its SRD value.

222 Taken altogether, the examples highlighted the opportunity to explore associations annotated
223 with low SRD values in mGWAS, beyond those with stringent significance cut-offs. Identifying
224 these potential false negative hits illustrates how useful the SRD annotation can be in discovering
225 biologically relevant results which would be missed when only considering stringent p-value cut-
226 offs in mGWAS reporting.

227

228 **Case study: mGWAS in Sickle Cell Disease patients**

229 To exemplify how the SRD metric can be used to prioritize mGWAS results, we present a new
230 mGWAS analysis performed in sickle cell disease (SCD) patients of African or African-
231 American ancestry. While the cause of SCD has been known for over a century, the molecular
232 determinants of the severity of this blood disease remain unknown and are influenced by genetic
233 variants unlinked to the beta-globin gene [31]. An mGWAS was performed in 651 SCD patients,
234 quantifying a total of 128 metabolites, to identify metabolites associated with genetic markers. A
235 total of 165 unique SNPs (with MAF > 1%) passed a genome-wide significance cut-off of $p <$
236 7.8125×10^{-10} (Supplementary Figure 2, 1×10^{-7} / 128 metabolites) and an additional unique 256
237 SNPs passed a suggestive cut-off of $p < 1 \times 10^{-7}$ (Supplementary Figure 2C). We identified a
238 total of eight loci associated with metabolite levels (Supplementary Figure 3), with four of them
239 reported in previous studies (Supplementary Table 2). The four additional associations were not
240 previously reported neither in the TK/HMC studies, nor in a large mGWAS of human blood
241 metabolites [32]. In all four cases, the frequency of the top associated SNP is larger in individuals
242 of African descent than in Europeans according to the gnomAD Genomes database, but these hits
243 were not found in the three largest mGWAS done in individuals of African ancestry to date [33-
244 35]. We recognize that while interesting, these novel associations will need replication in an
245 independent cohort.

246 We generated a QQplot based on minimum p-values for the gene-metabolite pairs extracted from
247 the SCD mGWAS (STAR Methods) and defined the gene-metabolite significance cut-off at $p <$
248 3.16×10^{-5} (Supplementary Figure 2D). We observed a significant negative correlation between
249 the mGWAS p-values and SRD values ($R = -0.1$, $p = 0.036$), replicating the result observed in

250 the TK study, demonstrating that this relationship between significance and SRD is reproducible
251 in a disease cohort, where biological mechanisms can be altered (Figure 3C). Given these
252 observations, we investigated whether associations above the gene-metabolite significance cut-
253 off of $p < 3.16 \times 10^{-5}$, but below genome-wide significant cut-offs usually required to report
254 associations, could be prioritized according to SRD values (candidate pairs below the plain line
255 in Figure 4). We split the associations into two suggestively significant categories, one that
256 includes association p-values between the genome-wide significance cut-off and $p < 1 \times 10^{-7}$
257 (Supplementary Figure 2C) and another that includes associations p-values between $p < 1 \times 10^{-7}$
258 and the gene-metabolite significance cut-off of $p < 3.16 \times 10^{-5}$ (Supplementary Figure 2D), which
259 we labelled S+ and S-, respectively. To evaluate the potential of association results in each
260 significance categories, we aimed to compute the proportion of small SRD values. From the
261 distribution of all SRD values computed on the KEGG overview graph (Methods), we observed
262 that 25% of SRD values (first quartile) are lower or equal to 8 (Supplementary Figure 4). We
263 thus used the SRD value ≤ 8 as a threshold for considering SRD values as small, reflecting close
264 biological relationships based on the graph topology. Furthermore, 90% of gene-metabolite
265 significant pairs in the well-powered TK and HMC studies have SRD ≤ 8 and, out of the genome-
266 wide significant associations with a numerical SRD value in this mGWAS (Block R, Figure 4),
267 two out of three have SRD ≤ 8 .

268 Two (40%) out of five associations from the S+ category have small SRD values, and 116/411 for
269 the S- category (28.22%). Next, we manually investigated the 118 suggestive associations with
270 SRD ≤ 8 by searching the literature for these gene-metabolite pairs and found six previously
271 described associations (5%): PRODH (proline dehydrogenase 1) and L-proline [11, 36, 37],
272 NT5C3A (5'-nucleotidase, cytosolic IIIA) and orotate [30], GATM (glycine amidinotransferase)
273 and creatine [38], UPB1 (beta-ureidopropionase 1) and 3-ureidopropionate [32], UGT2B17 (UDP
274 glucuronosyltransferase family 2 member B17) and sn-glycerol 3-phosphate (previously described
275 for total phosphoglycerides) [39], DGKH (diacylglycerol kinase eta) and tetradecanoyl-carnitine
276 [40], implying that these are likely to be real associations in SCD patients (Table 2). We thus
277 estimate that at least 5% of associations in this category are false negative and that the SRD metric
278 has added value in identifying them. It also demonstrates the ability of the SRD metric to retain
279 pairs that would not be otherwise reported, thus improving the mGWAS' potential for discovery.
280

281 **DISCUSSION**

282 The identification of an association between a SNP and a metabolite is usually supported solely
283 by the p-values of the statistical test without consideration for the known metabolic pathways.
284 Although the possibility of false positive hits is well understood among geneticists following up
285 on mGWAS results, the fact that reporting of gene-metabolite pairs using only statistical
286 significance can lead to an incomplete list of associations is less discussed. In this study, we
287 demonstrate how a simple metric called shortest reactional distance (SRD) can be useful for the
288 reporting and the *ad-hoc* annotation of gene-metabolite pairs for several p-values acceptance cut-
289 offs. By using previously published mGWAS, we have shown that the SRD is a metric capable of
290 retrospectively identifying a number of false negatives. The datasets we used in our study involved
291 different genomics and metabolomics protocols, different preprocessing strategies, different sets
292 of genetic variants (genotyped, imputed, exclusive to exons) and of metabolites, different
293 ethnicities (from European, African and Middle Eastern ancestries), disease-based and population
294 cohorts, illustrating the wide range of contexts in which our annotation is applicable.

295 PathQuant, the package developed to compute the SRD metric in this study, is not the only package
296 available that computes the shortest path metric between genes and metabolites. Similar to
297 PathQuant, MetaboSignal [13] is an R package based on the same mathematical criterion of
298 shortest path metric and also uses metabolic pathways from the KEGG database. There are,
299 however, important differences between these two methods. MetaboSignal combines both
300 metabolic and signaling pathway maps. While this can provide novel information about the
301 interaction between genes and metabolites that goes beyond the known enzymatic reactions and
302 pathways, it modifies the original topology of the curated pathways (signaling and metabolic).
303 These changes are not necessarily supported by biological data, which is crucial to ensure their
304 validity, as emphasized by Dumas et al. [17]. Thus, for our assessment of the potential of SRD
305 annotation in mGWAS, we favored a simpler approach that focuses only on curated metabolic
306 pathways, by converting a metabolic pathway map from KEGG into a graph of biochemical
307 reactions with metabolites as nodes and genes as edges. PathQuant computes SRD values from
308 any given list of gene-metabolite pairs using any given metabolic pathway graph in KGML format
309 thus keeping the original and curated topology of the metabolic pathways reported by KEGG. Of
310 note, a direct comparison of the shortest path metrics calculated by PathQuant and MetaboSignal
311 on the list of gene-metabolite pairs of the mGWAS investigated here could not be performed, as

312 the KEGG overview pathway graph we used with PathQuant is not accepted as input by
313 MetaboSignal. Indeed, MetaboSignal computes its shortest path metric using a custom graph that
314 is built from a pre-selected list of signaling and metabolic pathways.

315 The computation of the SRD metric with PathQuant within KEGG graphs leads to several possible
316 values: numerical, infinite or NA. Having a numerical annotation to a pair which is based on known
317 pathways leads to a better understanding of the underlying biology of an association. Indeed, we
318 have shown that SRD values decrease as statistical significance of gene-metabolite pairs increases:
319 this negative correlation between the level of significance and SRD values suggests an enrichment
320 of biologically relevant pairs as the p-value decreases, thus making SRD a promising annotation
321 metric to improve mGWAS reporting. We defined two categories for numerical values (short SRD:
322 $0 < \text{SRD} \leq 8$; large SRD: $\text{SRD} > 8$), which have been derived from the topological properties of
323 the KEGG overview graph (ID: hsa01100) as this graph contains all curated human biological
324 reactions, but different thresholds could be used in different contexts. The annotation of a pair with
325 an SRD lower or equal to 8 suggests that the gene, its expression levels or the protein it codes for,
326 may have a direct influence on the associated metabolite concentration. This is illustrated by the
327 statistical difference between the SRD values of mGWAS pairs and the null distribution of SRD
328 values within the KEGG overview graph and by manual investigation of specific pairs of interest,
329 such as the CPS1 and L-citrulline finding (SRD=1) in TK study, the UMPS and orotate example
330 (SRD=0) in HMC study. Moreover, for the SCD study, the gene-metabolite pairs presented in
331 Table 2 are annotated with an SRD lower or equal than 8 and have been already reported to be
332 associated with the corresponding metabolite levels in the literature. Despite the number of
333 investigated associations being very small for the R and S+ categories, we observed a decreasing
334 proportion of pairs with values lower than 8 and an increasing frequency of pairs with values
335 greater than 8, as p-values increases, in line with the negative correlation observed. These results
336 demonstrate, across all different mGWAS datasets we used, the potential of using the SRD metric
337 with a cut-off of $\text{SRD} \leq 8$ in order to identify false negatives hits and prioritize follow up studies
338 and experiments.

339 Another interesting case is when the SRD value is large, but the p-value is highly significant. In
340 this case, a direct influence of the gene on the associated metabolite is less clear. For example, we
341 noticed a large SRD value for the association between the alkaline phosphate (ALPL) and the
342 phosphocreatine (SRD = 22; p-value = 9.356×10^{-10}). One possibility is that this association could

343 have been a false negative. However, digging further into the literature, we found that
344 phosphocreatine can, in fact, be the substrate of the enzyme encoded by the ALPL gene [41, 42].
345 Thus, the true SRD value for ALPL-phosphocreatine pair should be 0, and the high SRD value
346 computed is a consequence of the current state of the publicly shared knowledge available on
347 KEGG. This result highlights a limitation of our approach, which is the incompleteness of some
348 metabolic pathways in KEGG. However, our SRD annotation pipeline of mGWAS results can help
349 identify those cases and could be useful to detect gaps and errors in metabolic network databases.
350 Other cases of $SRD > 8$ for highly significant association have been noted in our work, notably in
351 TK study (Figure 2A, Figure 2C; Figure 3A). When the path computed in KEGG is accurate and
352 the association is independently replicated, higher values of SRD can identify metabolic pathways
353 involving more indirect gene-metabolite relationships. By themselves, these cases are of interest
354 as they illustrate a violation of the hypothesis that biological proximity between a gene and a
355 metabolite is needed to regulate its level.

356 An SRD annotation with an infinite or NA value means that there was no path between the gene
357 and the metabolite within the chosen graph, because they are not connected (infinite value) or
358 because one of the two entities (or both) are missing from the queried graph (NA value). In most
359 cases, there likely exist no biological paths between these entities, but if the association is highly
360 significant, it could reflect actual gaps within the graph again. Investigation of enzymes or
361 metabolites with an unusual number of infinite or NA values may lead to more complete metabolic
362 pathway databases. Beyond these potential gaps in KEGG, other limitations of this resource are
363 the metabolic reactions provided do not specify cofactors, enzymatic complexes required for the
364 reaction to happen, and directionality, resulting in a decreasing level of complexity and accuracy
365 of the metabolic pathways. Thus, computing SRD on other resources for pathway mapping, such
366 as Recon3D, could improve our distance-based annotation pipeline. Indeed, this resource is
367 considered the most complete reconstruction of human metabolism so far [43, 44]. Furthermore,
368 in contrast to KEGG, Recon3D has the advantage of including more genes encoding transporters,
369 which would allow to increase the breath of SRD computation beyond enzymatic reactions. It also
370 includes information about cellular compartments and the directionality of reactions is known, in
371 contrast to KEGG graphs. Despite these advantages, the available file format (sbml) used in
372 Recon3D to represent metabolic pathways is incompatible with the graph structure we used for
373 KEGG. Given that the goal of this study was to establish if the SRD annotation is useful for

374 mGWAS annotation, a simple representation of metabolic pathways is highly appropriate and it is
375 outside the scope of the present study to implement an SRD metric based on Recon3D. Future
376 implementations on alternative databases should consider two additional desirable characteristics
377 of KEGG when comparing results : first, KEGG is the only database for which pathway maps were
378 built on known reactions from humans and others species, a clear advantage for non-human
379 mGWAS studies [45, 46]; second, KEGG pathway maps are built with an explicit labeling of side
380 compounds definition (such as ATP or H₂O) from KEGG RPAIRS [47], making their removal
381 from curated reactions possible, which eliminates shortcuts created by their over-representation
382 within the graph when computing SRD [48].

383 A crucial step for the annotation of gene-metabolite pairs with SRD values is to obtain the KEGG
384 IDs for the genes and metabolites. Although gene names and IDs are highly standardized [49] there
385 is a lack of uniformed nomenclature for metabolites. In most cases, there are multiple ways to refer
386 to a metabolite, with different studies using different reporting conventions, which complicated
387 the re-analysis of published datasets. For HMC and SCD studies, HMDB IDs [3] were provided
388 by the authors, which could be easily converted to KEGG IDs using the MetaboAnalyst Convert
389 tool. In the case of the TK study, we performed a manual annotation of the common names to
390 KEGG IDs, but this process was time consuming and required having the relevant expertise. This
391 manual work allowed us to include listed metabolites that did not have KEGG IDs, such as
392 acylcarnitines (see Material and Methods). These metabolites are measured in plasma but arise
393 from intracellular metabolism of their acyl-CoA counterparts, which in contrast to acylcarnitines
394 do not cross the plasma membrane. Based on the known direct link between the acylcarnitines and
395 the acyl-CoAs we have used the CoA counterparts to represent these metabolites, which do have
396 KEGG IDs. The increasing quantity of released studies involving metabolomics in the literature
397 encourages the community to release new standards to refer metabolites, such as Simplified
398 Molecular-Input LineEntry System [50], COordination of Standards in MetabOloromicS [51],
399 IUPAC International Chemical Identifier [52]. These efforts for standardisation of metabolites
400 naming, and the reporting improvements they allow, are leading to a promising leap for the field,
401 as this will result into better links being made between novel findings and already published data.
402 In practice, SRD annotations can be a great addition to bioinformatics pipelines in order to reduce
403 the number of potential candidate pairs to follow up on, by prioritizing hits based on a curated
404 biological information, as the number of released mGWAS grows. Furthermore, this metric could

405 be added within already available user-friendly resources such as mGWAS Explorer [14]. SRD
406 values could also play a role in designing targeted studies, in a context where extensive mGWAS
407 cannot be done or is not relevant, for example if the research question is about finding the genes
408 or proteins that regulate the levels of a specific metabolite, or conversely, finding which
409 metabolites are regulated by a specific enzyme. In this latter case, it would make sense to focus
410 more specifically on the metabolites showing low SRD values with the targeted enzyme. In the
411 future, we believe that extending the implementation by using other appropriate resources for
412 pathway mapping would likely improve coverage and enhance the value of the SRD metric in
413 mGWAS annotation pipelines.

414 In conclusion, we consider this work as a proof of concept of the benefits of using shortest
415 reactional distance metrics for annotating mGWAS results based on a model representation of the
416 human metabolism provided by the KEGG database. These metrics can also be used as a tool for
417 metabolic databases in order to more easily identify gaps within current metabolic pathway graphs
418 by using a new information provided by mGWAS. In this multi-omics era, we anticipate that large
419 scale studies looking for associations between genomics, proteomics and metabolomics signals
420 will soon become a new standard, as illustrated by recent studies in mice and humans [53, 54],
421 resulting in additional protein-metabolite pairs, for which the computation of SRD values may add
422 great value.

423

424 **STAR METHODS**

425 **RESOURCE AVAILABILITY**

426 **Lead contact**

427 Further information and requests should be directed to and will be fulfilled by the lead contact,
428 Julie Hussin (julie.hussin@umontreal.ca).

429

430 **Materials availability**

431 This study did not generate new materials.

432 **Data and code availability**

433 • Source data statement. This paper analyzed existing, publicly available summary mGWAS data
434 for the TK and HMC studies. The SCD study data comes from genomic and metabolomics of two
435 previously published studies. Raw data can be made available upon request to GL. Summary
436 statistics of SCD study will be made available publicly upon journal publication and from the lead
437 contact upon request in the meantime.

438 • All additional files necessary to reproduce the reported results of this paper are available at:
439 www.github.com/HussinLab/PathQuant/Publication/References.
440 • PathQuant source code is available at www.github.com/HussinLab/PathQuant/.
441 • Any additional information required to reanalyze the data reported in this paper is available from
442 the lead contact upon request.

443 **METHOD DETAILS**

444 **Overview of the mGWAS Datasets**

445 In this study we used three different mGWAS datasets. The TK study refers to an mGWAS [11]
446 previously performed on 7,824 participants (plasma or serum samples) from the TwinsUK cohort
447 [19] and the KORA study [20]. The mGWAS was carried out on 2.1 million genotyped SNPs
448 and 529 metabolites (assessed by targeted and untargeted metabolomics). The authors reported
449 299 SNP-metabolite (or ratio of metabolites) significant associations (at cut-off of $p \leq 1.03$

450 $\times 10^{-10}$ for metabolite concentrations and $p \leq 5.08 \times 10^{-13}$ for pairwise metabolite ratios) involving
451 187 unique metabolites and 145 loci, annotated as 132 causal genes [11]. The TK study also made
452 available mGWAS output files from METAL software [55] involving 486 metabolites without
453 any p-value cut-off [www.metabolomics.helmholtz-muenchen.de].

454 The HMC study refers to an mGWAS [21] previously performed on 614 Qatari participants
455 (plasma samples) from the Hamad Medical Corporation. The mGWAS was carried out on 1.6
456 million imputed exome variants and 826 metabolites (assessed by targeted and untargeted
457 metabolomics) and reported 3,127 significant associations (at cut-off of $p \leq 2.2 \times 10^{-10}$ for both
458 metabolite concentrations and pairwise metabolite ratios) in 21 locus-metabolite pairs. The
459 suggestive association results were also provided to the community, reporting all associations
460 with cut-off of $p \leq 1.4 \times 10^{-7}$ (Supplementary Data 7 in [21]), which include 6517 SNP-
461 metabolite (or ratio of metabolites) pairs with available p-values.

462 In the SCD study, an mGWAS was performed here using genetics and metabolomics data
463 published in different studies [22, 23] on a total of 651 SCD patients (plasma samples) including
464 401 individuals of African ancestry in the Genetic Modifier (GEN-MOD) cohort and 250
465 African-American individuals from Southwest USA in the Duke University Outcome Modifying
466 Genes (OMG) cohort. Metabolomics profiling was performed at the Broad Institute, and
467 appropriate statistical modelling was used to account for residual batch effects [56]. Briefly, for
468 association testing, 128 metabolites were profiled using a targeted approach in 651 plasma
469 samples from SCD patients. Metabolites were inverse normal transformed, adjusting for age and
470 sex. We then generated summary statistics for each cohort individually in a linear regression
471 model, accounting for relatedness using a kinship matrix as implemented in *rvtest* (v. 20171009)
472 [57] in GEN-MOD. The software *SNPTEST* [58] was employed in OMG to generate cohort-
473 specific summary statistics. We then meta-analyzed the effect size estimates and standard errors
474 from GEN-MOD and OMG using METAL [55]. After the pre-processing step, 6 million SNPs
475 were kept.

476

477 **Mapping shortest reactional distances (SRD) onto KEGG**

478 To compute the SRD metric, the R package PathQuant was developed [available at:
479 www.github.com/HussinLab/PathQuant]. PathQuant converts a metabolic pathway map into a

480 graph of biochemical reactions with metabolites as nodes and genes as edges (Figure 1). Briefly,
481 PathQuant takes as input a list of gene-metabolite pairs as pairs of KEGG identifiers (IDs) and a
482 list of metabolic pathways (eg. “hsa01100”, referred herein as KEGG overview graph, a global
483 concatenation of multiple distinct pathways). PathQuant then uses a KEGG XML file format
484 (KEGG Markup Language, KGML), downloaded using the KEGG API
485 [www.kegg.jp/kegg/rest/keggapi.html] to build the most up-to-date KEGG undirected graphs.
486 Next, PathQuant computes the SRD path between a gene and a metabolite from a given pair. The
487 SRD values are obtained using the breadth-first search Dijkstra algorithm [59]. PathQuant
488 outputs a text file containing genes and metabolites classification, Enzyme Commission number
489 (EC), KEGG Brite IDs, KEGG IDs of metabolic pathways for the SRD computation, and the
490 SRD values for all pairs. PathQuant also allows the visualization of SRD values annotation in a
491 heatmap, leading to a better identification of potentially interesting hits.

492

493 **Extracting and annotating gene-metabolite pairs from mGWAS summary statistics**

494 Using *bedtools* version v2.30.0 [60], we annotated each SNP to genes with KEGG IDs, using a
495 custom bed file: we downloaded gene coordinates for human genome build GRCh37.p13 from
496 NCBI (NCBI Homo sapiens Updated Annotation Release 105.2020/10/22, gff format) and
497 modified it to add 10,000 base pairs upstream and downstream for each of the 26,105 genes with
498 KEGG IDs. Because we focused on pairs involving genes encoding enzymes, which are the most
499 represented genes in the KEGG database, we only annotated SNPs with KEGG IDs classified as
500 enzymes within the KEGG Brite database by using the Brite enzyme code “BR:hsa01000”, leading
501 to a subset of 4,049/26,105 genes. For metabolites, if available, we used the provided IDs, either
502 KEGG IDs directly or HMDB IDs. For metabolites with HMDB ID the list was queried to
503 *Metaboanalyst* [61][www.metaboanalyst.ca/faces/upload/ConvertView.xhtml] to get the
504 corresponding KEGG IDs. For metabolites without KEGG or HMDB IDs provided, KEGG IDs
505 retrieval was achieved automatically by parsing the common names into the most standard KEGG
506 name format, and then queried the metabolite names with new format to MetaboAnalyst.
507 Additionally, specific metabolites were manually treated: acylcarnitines without KEGG IDs were
508 swapped to their acyl-CoA counterparts when available in KEGG. For each gene-metabolite pair
509 within a dataset (TK, HMC and SCD study), only one mGWAS p-value was kept, which is the
510 minimum p-value obtained for the association between any SNP annotated to the gene and that

511 specific metabolite. Based on that p-value, each pair is annotated as genome-wide significant,
512 suggestive, or non-significant: genome-wide p-value cut-offs were derived from a standard GWAS
513 Bonferroni correction approach (dividing by the number of tested metabolites). Suggestive
514 significance cut-offs, were obtained graphically from quantile-quantile plots (QQplots) of
515 minimum p-values for the variants-metabolite pairs of corresponding study. Specifically, we
516 determined the value on the QQplot x-axis where the slope starts to increase drastically.

517

518 **SRD null distribution and investigation of candidate gene-metabolite pairs**

519 To obtain a null distribution of SRD values from the KEGG overview graph (hsa01100), we
520 gathered all different KEGG IDs for both metabolites and genes found within the graph, and then
521 ran PathQuant on all possible pairs (Supplementary File 1). The KEGG overview graph version
522 (KGML v0.7.2 file) includes 1,351 genes and 2,889 metabolites, generating a total of 3,903,039
523 pairs for which we computed the SRD values, referred herein as the KEGG overview graph's SRD
524 null distribution. The first quartile of this distribution is used as the threshold to categorize any
525 pair with a close or far biological relationship label. This threshold was determined because 25%
526 of the smallest values are below half of the mean, leaving 75% as a representative sample of distant
527 relationships between genes and metabolites. We also performed a manual investigation of
528 candidate gene-metabolite pairs (Table 1, Supplementary Table 1) by extracting the corresponding
529 rsID using dbSNP annotation of UCSC Genome Browser [62], searching for published
530 associations of the metabolite with (1) the SNP rsID within GWAS Catalogue and/or with (2) the
531 gene symbol within PhenoScanner [6] (with parameters: cut-off p-value=1e-5, cut-off $r^2=0.8$,
532 build=37).

533

534 **Acknowledgements**

535 We thank all SCD participants for their contribution to this project. We thank members of the
536 Hussin group for helpful discussions and advice throughout this project, specifically Jean-
537 Christophe Grenier. We thank Fabien Jourdan, Clément Frainy and the iGenoMed Consortium for
538 initial discussions on this project. This work was completed thanks to computational resources
539 provided by Calcul Quebec clusters Narval, Beluga and Cedar, maintained by the Digital Research
540 Alliance of Canada. CB was supported by a PhD Scholarship from the NSERC-CREATE
541 Metabolomics Advanced Training and International Exchange (MATRIX) program. MR is a
542 FRQS Junior 1 research scholar, JGH is FRQS Junior 2 research scholar, GL holds the Canada
543 Research Chair (tier 1) in the genetics of heart and blood diseases, JDR holds a Canada Research
544 Chair in Genetics and Genomic Medicine. This study was supported by funding from the Montreal
545 Heart Institute Foundation and IVADO PRF Grant (PRF-2019-3378524797) to JGH and MR, and
546 a National Sciences and Engineering Research Council (NSERC) Discovery Grant (RGPIN-2022-
547 04262) to JGH, financial support of Génome Québec, Genome Canada, the government of Canada,
548 and the *Ministère de l'enseignement supérieur, de la recherche, de la science et de la technologie*
549 *du Québec*, the Canadian Institutes of Health Research (with contributions from the Institute of
550 Infection and Immunity, the Institute of Genetics, and the Institute of Nutrition, Metabolism and
551 Diabetes), Genome BC and Crohn's Colitis Canada (via Genome Canada grant number GPH-
552 129341) to CDR and JDR, Agilent Technologies Research Grant (3893 & 4075) to CDR. The
553 work in the GEN-MOD cohort is funded by the Canadian Institutes of Health Research (CIHR,
554 PJT #156248) and the Doris Duke Charitable Foundation. GEN-MOD sample and data collection
555 were supported by NIH grant HL-68922. The OMG-SCD study has been supported by grant
556 awards 2015131 and 2012126 from the Doris Duke Charitable Foundation, R01HL68959 and
557 R01HL079915 from the National Heart, Lung and Blood Institute, and DK110104 and DK124836
558 from the National Institute of Diabetes and Digestive and Kidney Diseases.

559 **Author contributions**

560 CB contributed to PathQuant package implementation with validation and support from RP and
561 PM, supervised by JGH. SC and STL developed and implemented an earlier version of the
562 PathQuant package and performed preliminary analyses of the TK study, supervised by MR, JDR,
563 GL and CDR. CB performed pre-processing of all mGWAS result datasets, generated final results

564 using PathQuant, interpreted the results and generated all figures and tables, supervised by CDR,
565 MR and JGH. GL, MJT, PB and AEAK generated the data from the SCD study. YB and MEG
566 performed the mGWAS of the SCD cohort, supervised by GL. CB and JGH wrote the manuscript,
567 with major input from MR and CDR. All authors revised and approved the final version of the
568 manuscript.

569 **Declaration of interests**

570 JGH received speaker honoraria from Dalcor and District 3 Innovation Centre. PB is consultant
571 for ADDMEDICA, NOVARTIS, ROCHE, GBT, Bluebird, EMMAUS, HEMANEXT, AGIOS,
572 VERTEX, received lecture fees for NOVARTIS, ADDMEDICA, AGIOS, JAZZPHARMA,
573 VERTEX, is on steering committee for NOVARTIS and ADDMEDICA, receives research support
574 from ADDMEDICA, foundation Fabre, NOVARTIS, Bluebird, EMMAUS, GBT, and is
575 cofounder of INNOVHEM.

REFERENCES:

1. Gallagher, M.D. and A.S. Chen-Plotkin, *The Post-GWAS Era: From Association to Function*. Am J Hum Genet, 2018. **102**(5): p. 717-730.
2. Wishart, D.S., *Metabolomics for Investigating Physiological and Pathophysiological Processes*. Physiol Rev, 2019. **99**(4): p. 1819-1875.
3. Wishart, D.S., et al., *HMDB: the Human Metabolome Database*. Nucleic Acids Research, 2007. **35**(suppl_1): p. D521-D526.
4. Pe'er, I., et al., *Estimation of the multiple testing burden for genomewide association studies of nearly all common variants*. Genet Epidemiol, 2008. **32**(4): p. 381-5.
5. Buniello, A., et al., *The NHGRI-EBI GWAS Catalog of published genome-wide association studies, targeted arrays and summary statistics 2019*. Nucleic Acids Res, 2019. **47**(D1): p. D1005-D1012.
6. Kamat, M.A., et al., *PhenoScanner V2: an expanded tool for searching human genotype-phenotype associations*. Bioinformatics, 2019. **35**(22): p. 4851-4853.
7. Elsworth, B., et al., *The MRC IEU OpenGWAS data infrastructure*. 2020, bioRxiv.
8. Ghousaini, M., et al., *Open Targets Genetics: systematic identification of trait-associated genes using large-scale genetics and functional genomics*. Nucleic Acids Res, 2021. **49**(D1): p. D1311-D1320.
9. Shashkova, T.I., et al., *PheLiGe: an interactive database of billions of human genotype-phenotype associations*. Nucleic Acids Res, 2021. **49**(D1): p. D1347-D1350.
10. Pinero, J., et al., *The DisGeNET knowledge platform for disease genomics: 2019 update*. Nucleic Acids Res, 2020. **48**(D1): p. D845-D855.
11. Shin, S.Y., et al., *An atlas of genetic influences on human blood metabolites*. Nat Genet, 2014. **46**(6): p. 543-550.
12. Gagliano Taliun, S.A., et al., *Exploring and visualizing large-scale genetic associations by using PheWeb*. Nat Genet, 2020. **52**(6): p. 550-552.
13. Rodriguez-Martinez, A., et al., *MetaboSignal: a network-based approach for topological analysis of metabotype regulation via metabolic and signaling pathways*. Bioinformatics, 2017. **33**(5): p. 773-775.
14. Chang, L., et al., *mGWAS-Explorer: Linking SNPs, Genes, Metabolites, and Diseases for Functional Insights*. Metabolites, 2022. **12**(6).
15. Kanehisa, M. and S. Goto, *KEGG: kyoto encyclopedia of genes and genomes*. Nucleic Acids Res, 2000. **28**(1): p. 27-30.
16. Pavlopoulos, G.A., et al., *Using graph theory to analyze biological networks*. BioData Min, 2011. **4**: p. 10.
17. Dumas, M.E., et al., *Topological analysis of metabolic networks integrating co-segregating transcriptomes and metabolomes in type 2 diabetic rat congenic series*. Genome Med, 2016. **8**(1): p. 101.
18. Fuhrer, T., et al., *Genomewide landscape of gene-metabolome associations in Escherichia coli*. Mol Syst Biol, 2017. **13**(1): p. 907.
19. Moayyeri, A., et al., *The UK Adult Twin Registry (TwinsUK Resource)*. Twin Res Hum Genet, 2013. **16**(1): p. 144-9.
20. Wichmann, H.E., et al., *KORA-gen--resource for population genetics, controls and a broad spectrum of disease phenotypes*. Gesundheitswesen, 2005. **67** Suppl 1: p. S26-30.

21. Yousri, N.A., et al., *Whole-exome sequencing identifies common and rare variant metabolic QTLs in a Middle Eastern population*. Nat Commun, 2018. **9**(1): p. 333.
22. Pincez, T., et al., *Variation and impact of polygenic hematological traits in monogenic sickle cell disease*. Haematologica, 2022.
23. Ilboudo, Y., et al., *Potential causal role of l-glutamine in sickle cell disease painful crises: A Mendelian randomization analysis*. Blood Cells Mol Dis, 2021. **86**: p. 102504.
24. Martin, A.R., et al., *Clinical use of current polygenic risk scores may exacerbate health disparities*. Nature Genetics, 2019. **51**(4): p. 584-591.
25. Morales, J., et al., *A standardized framework for representation of ancestry data in genomics studies, with application to the NHGRI-EBI GWAS Catalog*. Genome Biology, 2018. **19**(1): p. 21.
26. Chen, Z., et al., *Revisiting the genome-wide significance threshold for common variant GWAS*. G3 (Bethesda), 2021. **11**(2).
27. Marco-Marín, C., et al., *Δ1-Pyrroline-5-carboxylate synthetase deficiency: An emergent multifaceted urea cycle-related disorder*. 2020. **43**(4): p. 657-670.
28. Lotta, L.A., et al., *A cross-platform approach identifies genetic regulators of human metabolism and health*. Nat Genet, 2021. **53**(1): p. 54-64.
29. Hysi, P.G., et al., *Metabolome Genome-Wide Association Study Identifies 74 Novel Genomic Regions Influencing Plasma Metabolites Levels*. Metabolites, 2022. **12**(1).
30. Rhee, E.P., et al., *A genome-wide association study of the human metabolome in a community-based cohort*. Cell Metab, 2013. **18**(1): p. 130-43.
31. Kato, G.J., et al., *Sickle cell disease*. Nature Reviews Disease Primers, 2018. **4**(1): p. 18010.
32. Long, T., et al., *Whole-genome sequencing identifies common-to-rare variants associated with human blood metabolites*. Nat Genet, 2017. **49**(4): p. 568-578.
33. Liu, C., et al., *Crosstalk between Host Genome and Metabolome among People with HIV in South Africa*. Metabolites, 2022. **12**(7).
34. Yu, B., et al., *Genetic determinants influencing human serum metabolome among African Americans*. PLoS Genet, 2014. **10**(3): p. e1004212.
35. Luo, S., et al., *Genome-wide association study of serum metabolites in the African American Study of Kidney Disease and Hypertension*. Kidney Int, 2021. **100**(2): p. 430-439.
36. Demirkan, A., et al., *Insight in genome-wide association of metabolite quantitative traits by exome sequence analyses*. PLoS Genet, 2015. **11**(1): p. e1004835.
37. Draisma, H.H.M., et al., *Genome-wide association study identifies novel genetic variants contributing to variation in blood metabolite levels*. Nat Commun, 2015. **6**: p. 7208.
38. Barton, A.R., et al., *Whole-exome imputation within UK Biobank powers rare coding variant association and fine-mapping analyses*. Nat Genet, 2021. **53**(8): p. 1260-1269.
39. Kettunen, J., et al., *Genome-wide study for circulating metabolites identifies 62 loci and reveals novel systemic effects of LPA*. Nat Commun, 2016. **7**: p. 11122.
40. Suhre, K., et al., *Human metabolic individuality in biomedical and pharmaceutical research*. Nature, 2011. **477**(7362): p. 54-60.
41. Fedde, K.N. and M.P. Whyte, *Alkaline phosphatase (tissue-nonspecific isoenzyme) is a phosphoethanolamine and pyridoxal-5'-phosphate ectophosphatase: normal and hypophosphatasia fibroblast study*. Am J Hum Genet, 1990. **47**(5): p. 767-75.

42. Di Mauro, S., et al., *Kinetic characterization of hypophosphatasia mutations with physiological substrates*. J Bone Miner Res, 2002. **17**(8): p. 1383-91.
43. Preciat Gonzalez, G.A., et al., *Comparative evaluation of atom mapping algorithms for balanced metabolic reactions: application to Recon 3D*. J Cheminform, 2017. **9**(1): p. 39.
44. Brunk, E., et al., *Recon3D enables a three-dimensional view of gene variation in human metabolism*. Nat Biotechnol, 2018. **36**(3): p. 272-281.
45. Chen, W., et al., *Genome-wide association analyses provide genetic and biochemical insights into natural variation in rice metabolism*. Nature Genetics, 2014. **46**(7): p. 714-721.
46. Chen, W., et al., *Comparative and parallel genome-wide association studies for metabolic and agronomic traits in cereals*. Nat Commun, 2016. **7**: p. 12767.
47. Cottret, L. and F.J.P. Jourdan, *Graph methods for the investigation of metabolic networks in parasitology*. 2010. **137**(9): p. 1393-1407.
48. Faust, K., D. Croes, and J.J.J.o.m.b. Van Helden, *Metabolic pathfinding using RPAIR annotation*. 2009. **388**(2): p. 390-414.
49. Eyre, T.A., et al., *The HUGO Gene Nomenclature Database, 2006 updates*. Nucleic Acids Research, 2006. **34**(suppl_1): p. D319-D321.
50. Weininger, D.J.J.o.c.i. and c. sciences, *SMILES, a chemical language and information system. 1. Introduction to methodology and encoding rules*. 1988. **28**(1): p. 31-36.
51. Salek, R.M., et al., *COordination of Standards in MetabOlonicS (COSMOS): facilitating integrated metabolomics data access*. Metabolomics, 2015. **11**(6): p. 1587-1597.
52. Heller, S.R., et al., *InChI, the IUPAC International Chemical Identifier*. J Cheminform, 2015. **7**: p. 23.
53. Williams, E.G., et al., *Systems proteomics of liver mitochondria function*. Science, 2016. **352**(6291): p. aad0189.
54. Suhre, K. and S. Zaghloul, *Connecting the epigenome, metabolome and proteome for a deeper understanding of disease*. J Intern Med, 2021. **290**(3): p. 527-548.
55. Willer, C.J., Y. Li, and G.R. Abecasis, *METAL: fast and efficient meta-analysis of genomewide association scans*. Bioinformatics, 2010. **26**(17): p. 2190-1.
56. Leek, J.T., et al., *The sva package for removing batch effects and other unwanted variation in high-throughput experiments*. Bioinformatics, 2012. **28**(6): p. 882-3.
57. Zhan, X., et al., *RVTESTS: an efficient and comprehensive tool for rare variant association analysis using sequence data*. Bioinformatics, 2016. **32**(9): p. 1423-6.
58. Marchini, J., et al., *A new multipoint method for genome-wide association studies by imputation of genotypes*. Nat Genet, 2007. **39**(7): p. 906-13.
59. Dijkstra, E.W., *A note on two problems in connexion with graphs*. Numerische Mathematik, 1959. **1**(1): p. 269-271.
60. Quinlan, A.R. and I.M. Hall, *BEDTools: a flexible suite of utilities for comparing genomic features*. Bioinformatics, 2010. **26**(6): p. 841-2.
61. Pang, Z., et al., *Using MetaboAnalyst 5.0 for LC-HRMS spectra processing, multi-omics integration and covariate adjustment of global metabolomics data*. Nat Protoc, 2022. **17**(8): p. 1735-1761.
62. Kent, W.J., et al., *The human genome browser at UCSC*. 2002. **12**(6): p. 996-1006.

MAIN FIGURES AND LEGENDS

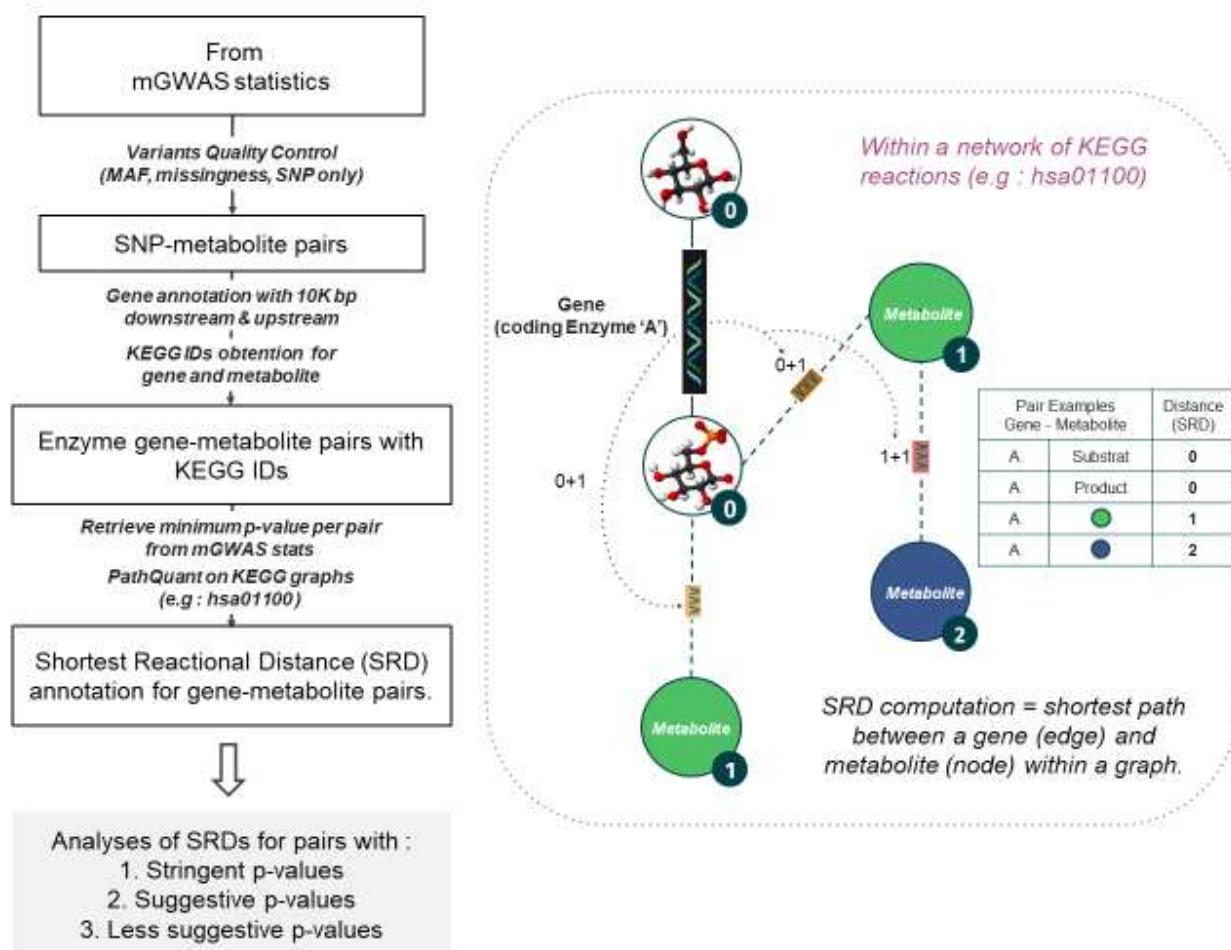


Figure 1. Overview of the Shortest Reactional Distance (SRD) annotation process.

On the left panel, we show the main steps of the pipeline used to annotate mGWAS datasets. Boxes indicate input and outputs of processing steps described on arrows. The right panel shows the meaning of an SRD annotation of a reaction. The queried metabolite is in blue, the queried gene is Enzyme A and a sub-graph is shown. The SRD between Enzyme A and its main reactants, the substrate and the product, are of 0, everything running deeper adds 1 at each step (dashed arrows): green metabolites are at SRD = 1 and the queried metabolite is at SRD = 2 from Enzyme A.

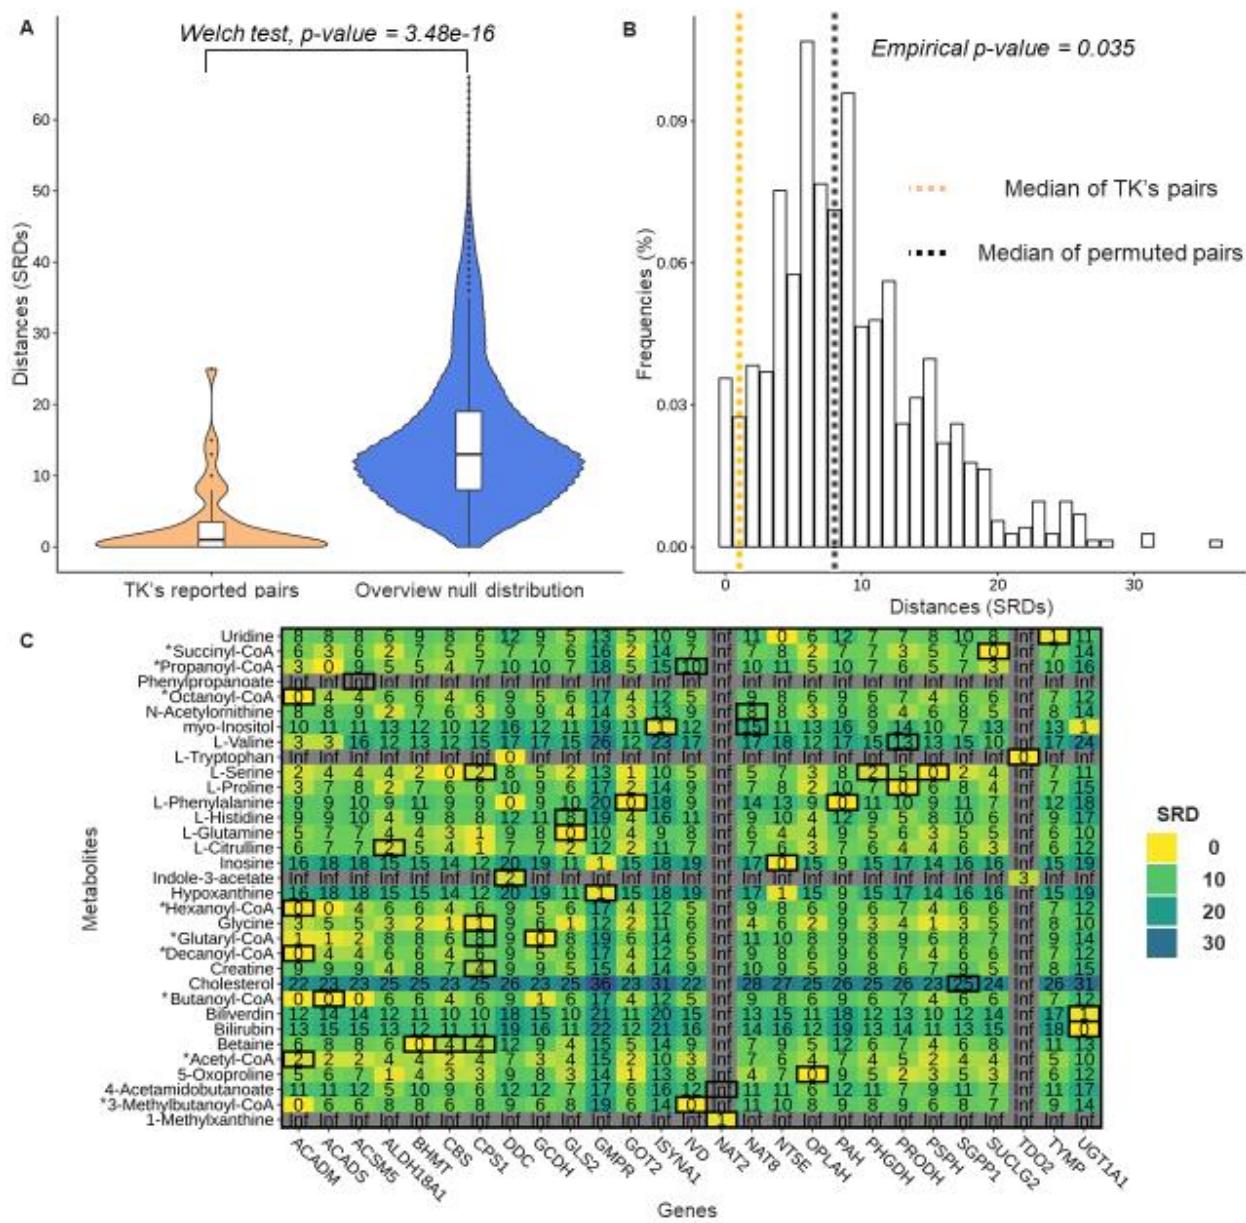


Figure 2. SRD of stringently associated gene-metabolite pairs in the TK study.

(A) Comparison of SRD annotations for reported genome-wide significant associations from the TK study (orange) and the distribution of all SRD values within KEGG overview graph (hsa01100) (B) Distribution of SRD values computed from permuted gene-metabolite pairs from TK study, with a median SRD of 8 (black dotted line). The median SRD of 1 (orange dotted line) represents an empirical p -value of $p=0.035$. (C) Heatmap representing all genes and metabolites included in reported genome-wide significant associations from the TK study. The 40 mapped gene-metabolite associations are enclosed in black boxes on the heatmap ($n=40$). Abbreviations: CoA = Coenzyme A.

* CoA metabolites are proxies for measured carnitines.

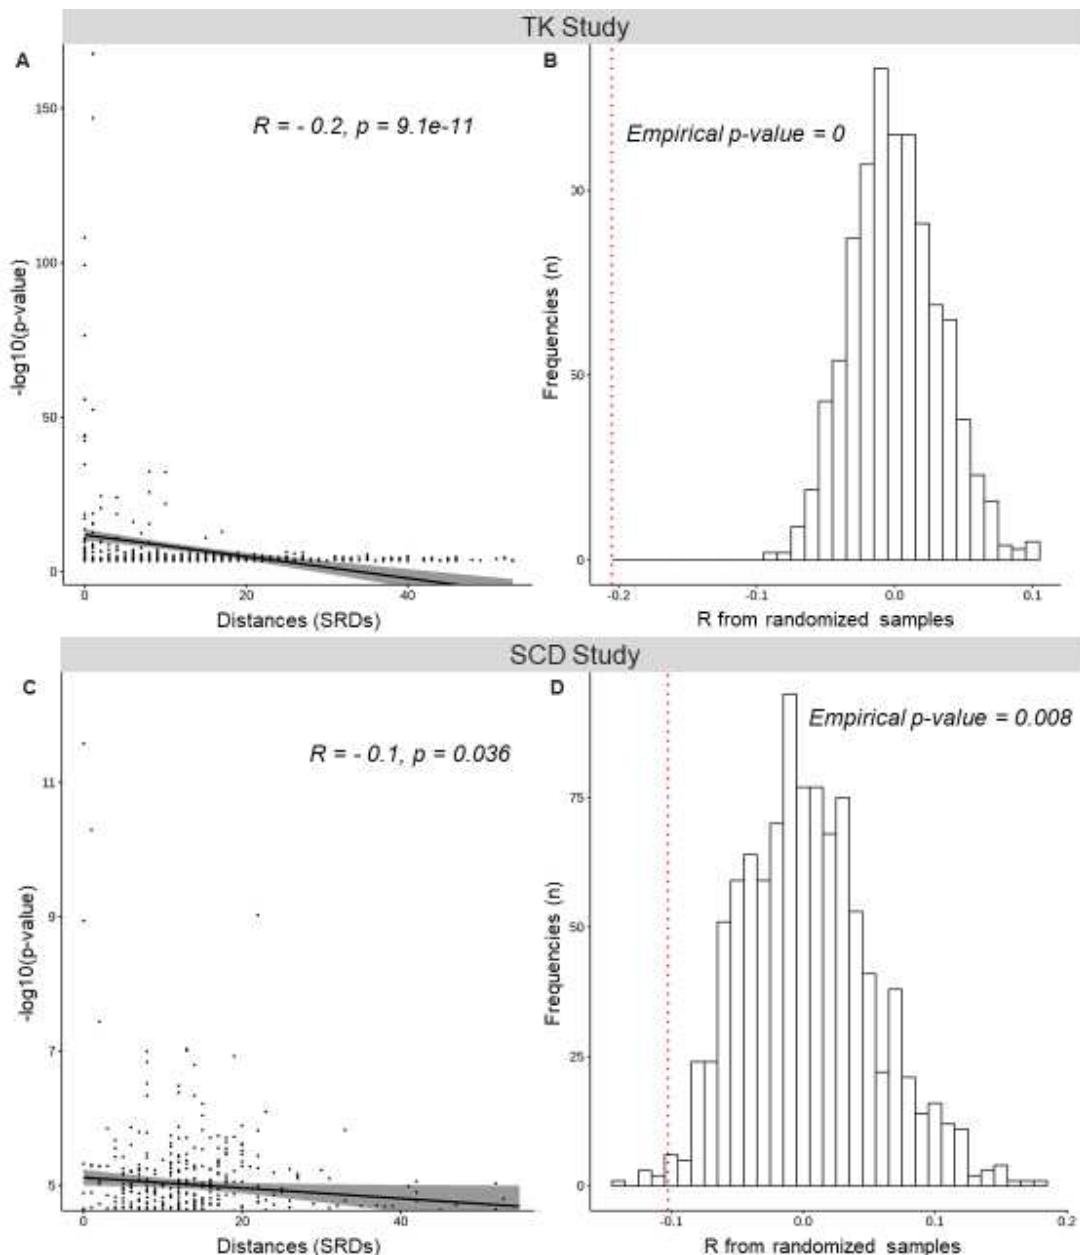


Figure 3. Relationship between p-values and SRD values for TK and SCD studies

Correlation plot between the $-\log_{10}(p\text{-value})$ and SRD values for gene-metabolite pairs in (A) the TK study (p-value cutoff was 3.16×10^{-4} , Supplementary Figure 3B) and (C) the SCD study (p-value cutoff was 3.16×10^{-5} , Supplementary Figure 3E). Correlation computed on permuted data ($N=1000$) to take into account the graph structure allowed computation of empirical p-values for TK (B) and SCD (C), with the true correlation coefficient presented (dotted red line).

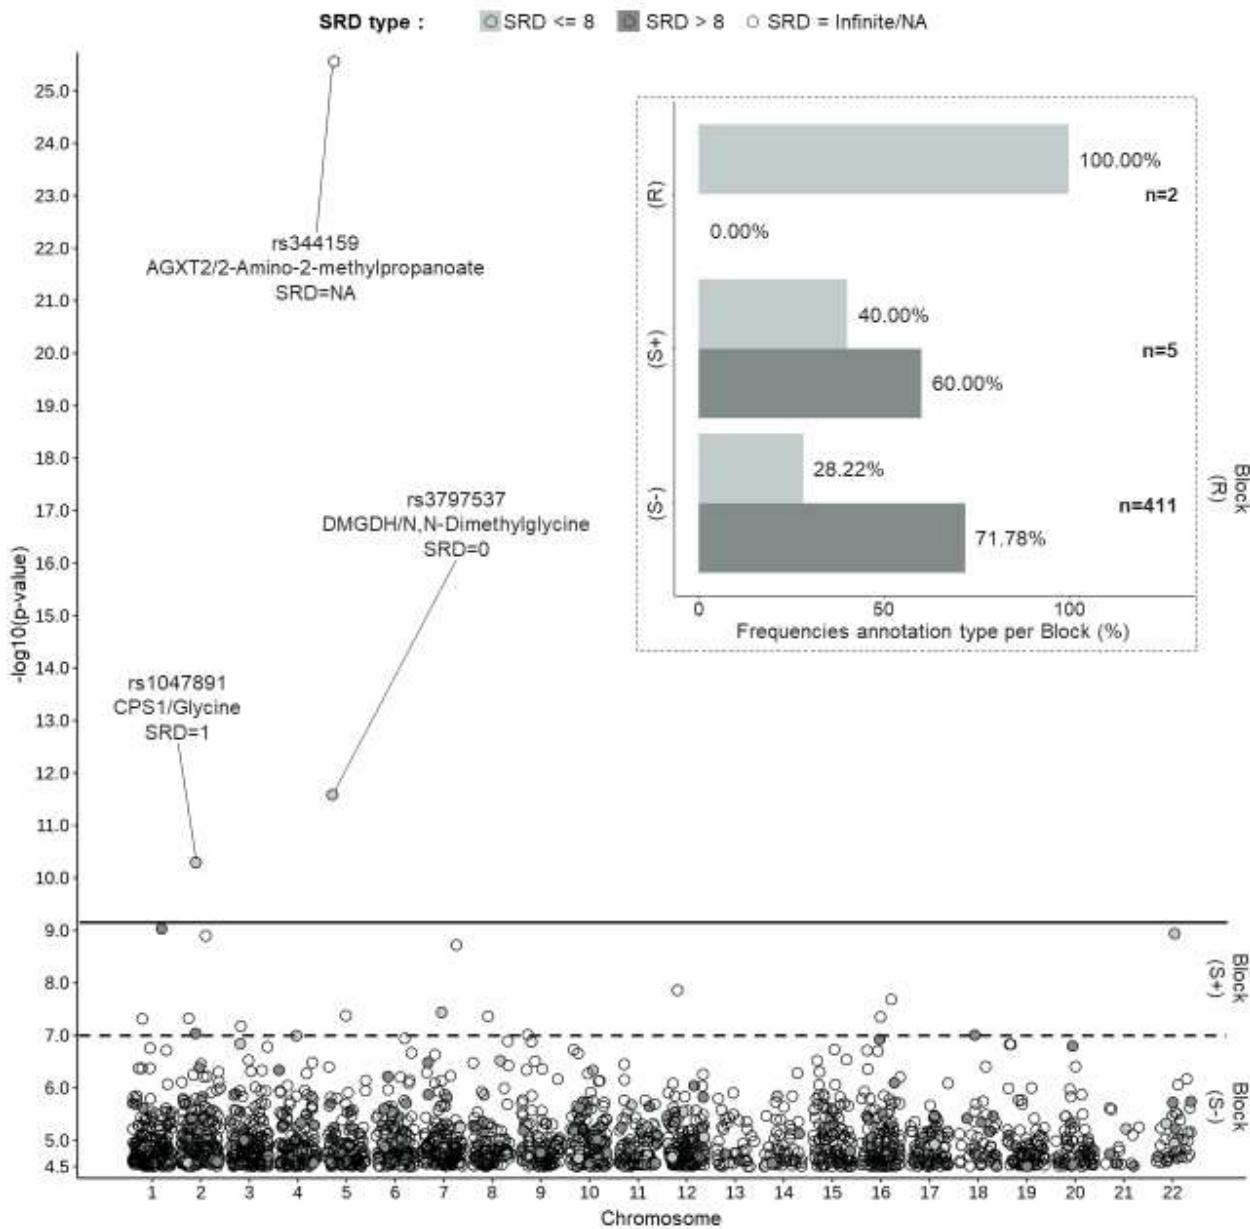


Figure 4. Identification of potential false negative hits for different p-values cut-offs in the SCD study. The main panel represents $-\log_{10}(p\text{-values})$ for gene-metabolite pairs for SNPs on each chromosome (x-axis). Coloring represents the SRD annotation from 3 categories: SRD = Infinite or NA (white), numerical SRD ≤ 8 (light grey), numerical SRD > 8 (dark grey). On the top right panel, the histogram represents the frequencies of SRD ≤ 8 and > 8 for each block. The frequency percentage is obtained by summing only the pairs with a numerical SRD value (n) within each category according to different p-value cut-offs: R : $p < 7.8125 \times 10^{-10}$, S+ : $7.8125 \times 10^{-10} < p < 1 \times 10^{-7}$ and S- : $1 \times 10^{-7} < p < 3.16 \times 10^{-5}$.

MAIN TABLES AND LEGENDS

Study name	TK	HMC	SCD	TK-None
Summary stats cut-off	1.03×10^{-10}	1.79×10^{-7}	None	None
Sample size	7,824	614	651	7,824
SNP-Metabolite name pairs (unique SNP - unique metabolites)	336 (218-186)	6 465 (3 192-100)	3 139 070 208 (24 523 986-128)	1 236 909 025 (2 617 408-486)
Enzyme Gene-Metabolite pairs with KEGG IDs: PathQuant input (unique genes - unique metabolites)	74 (43 - 56)	68 (32-20)	209 645 290 (3875-118)	46 779 684 (3815-177)
Pairs' SRD annotations within KEGG Overview Graph (unique genes - unique metabolites*)	SRD: 38 Inf: 2 NA: 34 (27 - 33)	SRD: 38 Inf: 4 NA: 26 (13 - 8)	SRD: 82 145 Inf: 38 980 NA: 336 125 (1275 - 95)	SRD: 86 219 Inf: 55 822 NA: 533 214 (1257 - 113)

Table 1. Dataset description of used dataset for the TK, HMC, and SCD studies.

Line 1: Dataset label for TK, HMC and SCD studies and the TK dataset without p-values cut-off (TK-None). Line 2: Maximum p-value cut-off. Line 3: Maximum number of samples available for the mGWAS study. Line 4: Number of SNP-metabolites pairs. Line 5: Number of gene coding for enzyme-metabolite pairs for which KEGG IDs have been obtained. Line 6: Number of SRD annotation within map hsa01100 (KEGG overview graph) categorized in three categories. Line 4,5,6 have the count of unique genes and metabolites involved in pairs between parenthesis. *: the number of unique genes and metabolites are ignoring NA values. Abbreviations: ID = Identifier; SNP = Single-nucleotide polymorphism; SRD = Shortest Reactional Distance; Inf = Infinite value; NA = Not available, a missing annotation.

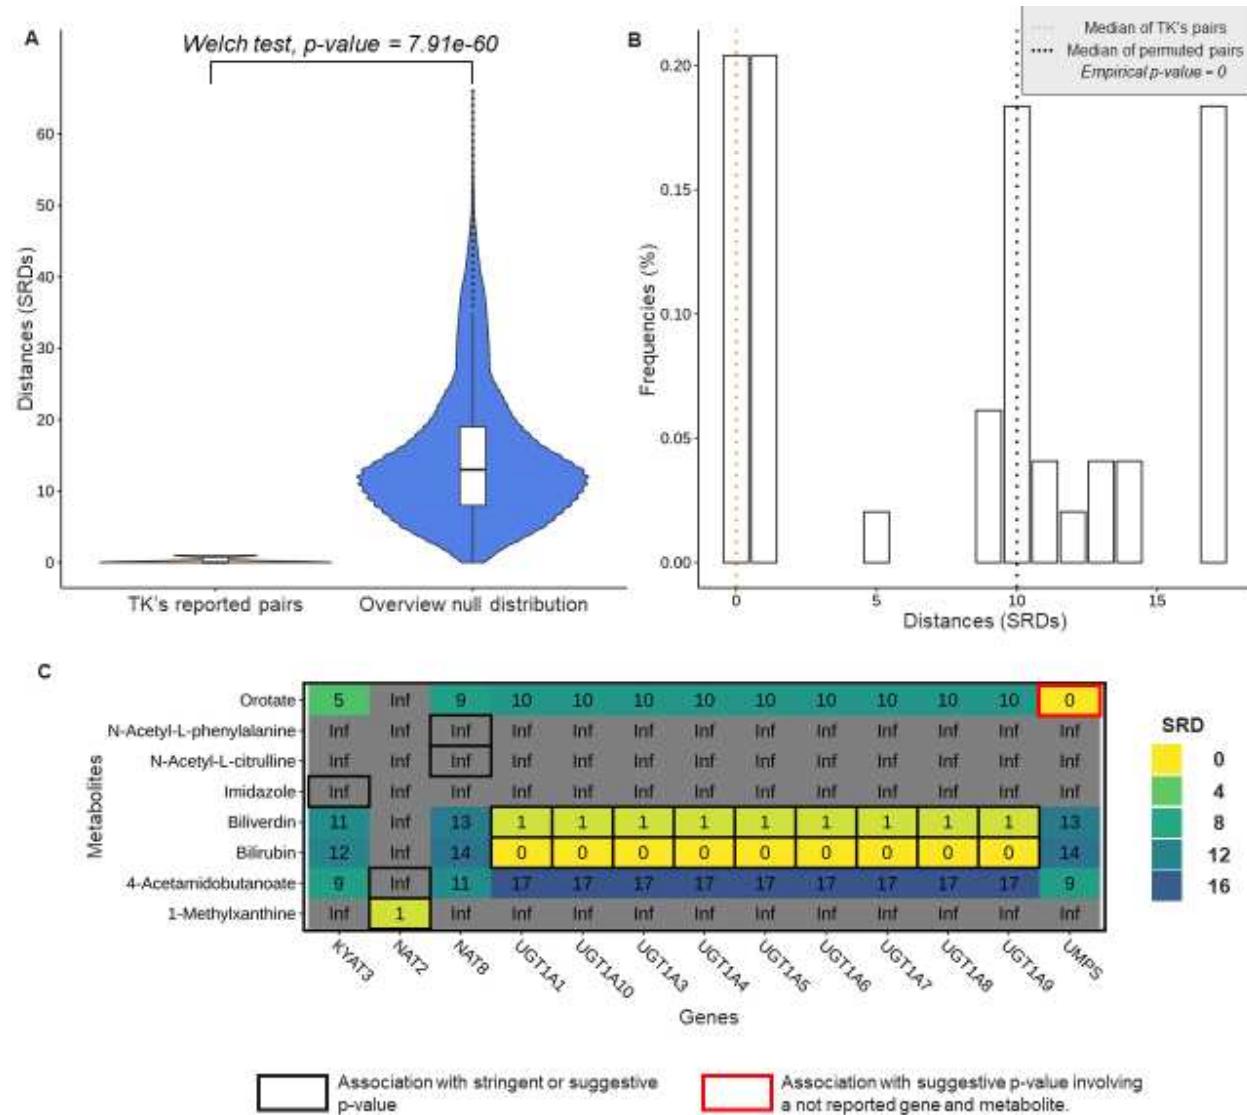
SNP (Chromosome_Position_Allele1_Allele2)	Gene symbol (KEGG ID)	Metabolite common name (KEGG ID)	P-value	SRD	Comments – Publication ID reporting the described association
22_18905964_C_T	PRODH (hsa:5625)	L-Proline (C00148)	4.77 4e-06	0	rs2904552, PMID:24816252, 26068415, 25569235 http://www.phenosscanner.medschl.cam.ac.uk/?query=PRODH&catalogue=mQTL&p=1e-5&proxies=None&r2=0.8&build=37
7_33081514_A_G	NT5C3A (hsa:51251)	Orotate (C00295)	3.65 9e-08	2	rs4316067, PMID: 23823483 http://www.phenosscanner.medschl.cam.ac.uk/?query=NT5C3A&catalogue=mQTL&p=1e-5&proxies=None&r2=0.8&build=37
15_45682944_T_C	GATM (hsa:2628)	Creatine (C00300)	5.16 0e-06	1	rs536148271, PMID : 34226706, http://www.phenosscanner.medschl.cam.ac.uk/?query=GATM+&catalogue=mQTL&p=1e-5&proxies=None&r2=0.8&build=37
22_24893867_G_A	UPB1 (hsa:51733)	3-Ureidopropionate (C02642)	1.14 9e-09	0	No rsID found, PMID:28263315, http://www.phenosscanner.medschl.cam.ac.uk/?query=UPB1&catalogue=mQTL&p=1e-5&proxies=None&r2=0.8&build=37
4_69444550_C_A	UGT2B17 (hsa:7367)	sn-Glycerol 3-phosphate (C00093)	2.93 e-05	6	No rsID found, PMID : 27005778 (Total phosphoglycerides) http://www.phenosscanner.medschl.cam.ac.uk/?query=UGT2B17&catalogue=mQTL&p=1e-5&proxies=None&r2=0.8&build=37
13_42683873_C_T	DGKH (hsa:160851)	Tetradecanoyl-CoA* (C02593)	2.72 e-05	8	rs75732304 PMID : 21886157 (2-tetradecenoyl carnitine/gamma-glutamylmethionine*) http://www.phenosscanner.medschl.cam.ac.uk/?query=DGKH&catalogue=mQTL&p=1e-5&proxies=None&r2=0.8&build=37

Table 2: Associations with SRD lower or equal to 8 for the SCD study, illustrating potential false negative hits discovered by manual verification.

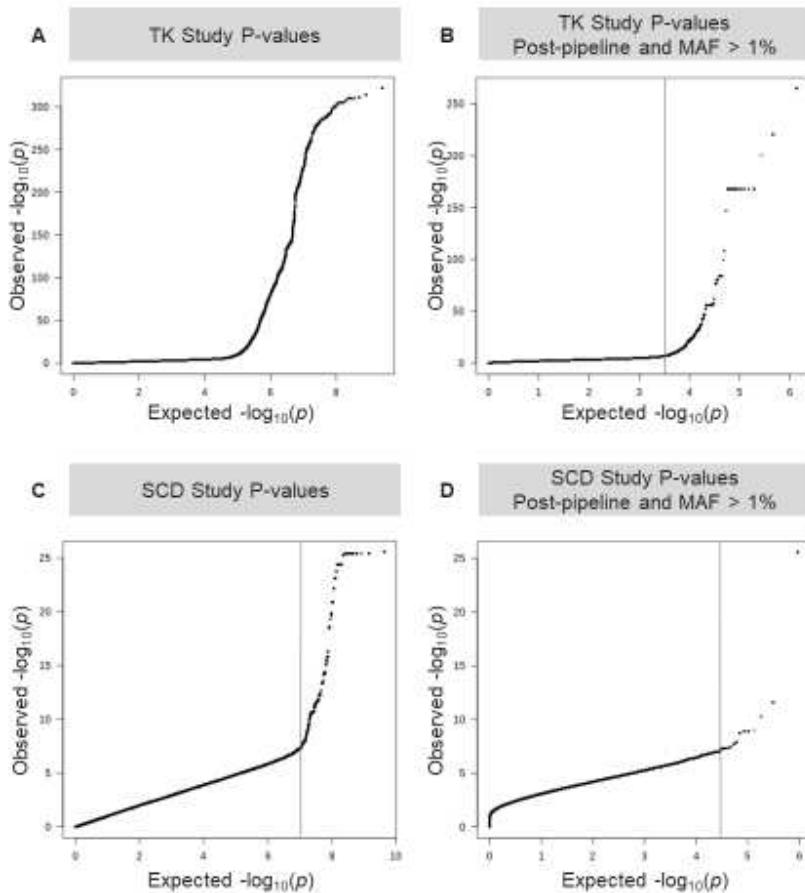
Column 1: SNP information formatted as CHROMOSOME_POSITION_ALLELE1_ALLELE2 in hg19. Column 2: Gene symbol found the in the KEGG ID entry. Column 3: Metabolite common name found the in the KEGG ID entry. Column 4: P-value of the association between the SNP and the metabolite. Column 5: SRD annotation of the gene-metabolite pair using KEGG graph hsa01100. Column 6: Miscellaneous information for the association: rsID if available; PMID of the publications reporting association between the gene, or the SNP with the associated metabolite; link to the query used in Phenoscanner website. Abbreviations: ID = Identifier; SNP = Single-nucleotide polymorphism; SRD = Shortest Reactional Distance, PMID = PubMed Identifier. * CoA are representatives of the measured carnitines.

SUPPLEMENTAL INFORMATION

Supplementary figures and legends.

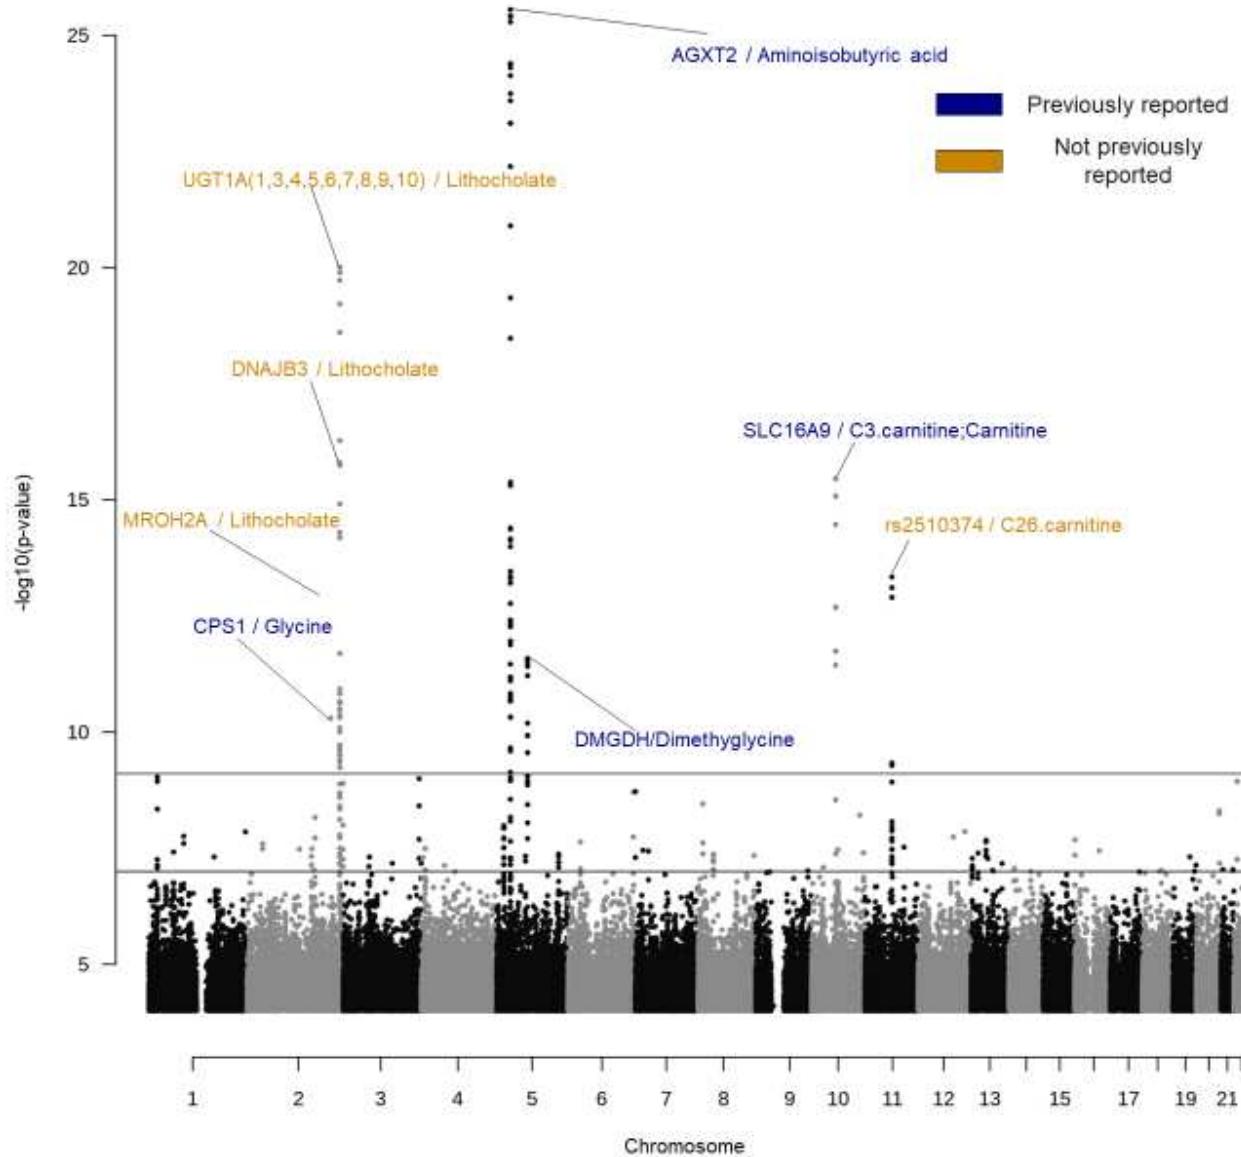


Supplementary Figure 1. SRD of stringently and suggestive associated gene-metabolite pairs in the HMC study. (A) Comparison of SRD annotations for reported genome-wide significant associations from the HMC study (orange) and the distribution of all SRD values within KEGG overview graph (hsa01100) (B) Distribution of SRD values computed from permuted gene-metabolite pairs from HMC study, with a median SRD of 10 (black dotted line). The median SRD of 0 (orange dotted line) represents an empirical p -value of $p=0$. (C) Heatmap representing all genes and metabolites included in reported stringent and suggestive associations from the HMC study. The 24 mapped gene-metabolite associations are enclosed: a black box indicates a reported stringent or a suggestive pair for which either the gene, the metabolite or both have been reported in the dataset ($n=23$) while a red enclosed box indicates a suggestive pair for which the gene and the metabolite have not been reported in the study ($n=1$).



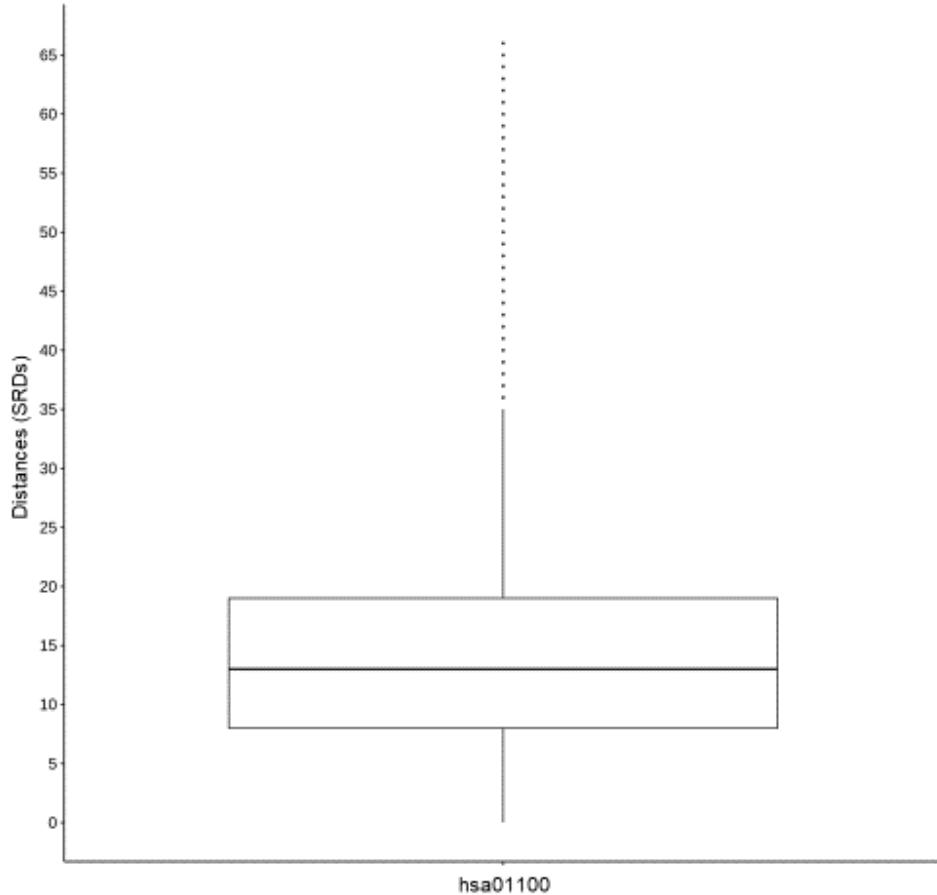
Supplementary Figure 2. QQplots for the TK and SCD studies.

(A/C) QQplots using all p-values from mGWAS files available, before any pre-processing steps for TK study (A) and for the SCD study (C). (B/D) QQplots using p-values at the end of the pipeline (See Methods) for the TK study (B) and for the SCD study (D). Vertical lines indicate the graphically determined cut-offs we used in our Results section. Abbreviations: MAF = Minor Allele Frequency, QQplot = Quantile-Quantile plot.



Supplementary Figure 3. Manhattan plot of the SCD study.

Manhattan plot of the SCD study with a p-values below 1×10^{-4} cutoff, each dot represents an association between a variant and a metabolite. Black and white is an alternating color palette for chromosomes. The first line from plot bottom (dark plain) indicates the genome-wide threshold at 1×10^{-7} , and the second line (light plain) at 7.8125×10^{-10} , indicates the first threshold corrected by Bonferroni (See Methods). Top hits for each mapped gene above the Bonferroni threshold are labelled with two distinct colors: blue if the association has not been previously reported, or orange if the association has been already reported in the literature. Gene symbols are separated to metabolite name(s) by "/" symbol.



Min.	1st Qu.	Median	Mean	3rd Qu.	Max.
0.0	8.0	13.0	14.9	19.0	66.0

Supplementary Figure 4. Summary statistics of the null distribution of SRD values within the KEGG overview graph.

Boxplot of the SRD values and table of summary statistics for the null distribution of SRD values within the KEGG overview graph (hsa01100). The mean SRD within this graph is 14.89, with a maximum value of 66. The first quartile (SRD <8) is used as the threshold to categorize any pair with a close or far biological relationship label.

SUPPLEMENTARY METHODS

PathQuant tool to compute shortest reactional distances (SRD).

The method was developed using R and delivered as an R package. PathQuant (for Pathway Quantity).

1.1 Input parameters

Gene-metabolite pairs: PathQuant accepts the following input for SRD computation: 1) gene(s); and 2) their associated metabolite(s) in columns, each with their specific KEGG identifiers (IDs). Each row represents a unique gene-metabolite association. Only associations between genes and individual metabolites are taken as entry, hence associations between a gene and a ratio of metabolites must be separated prior to analysis (usually by generating gene-numerator/denominator pairs).

Selected metabolic pathways: PathQuant accepts a list of metabolic pathways, each with its specific KEGG IDs. For this application, we used the map ID: hsa: 01100 in KEGG (referred herein as the ‘overview’ or ‘hsa01100’ throughout this article).

1.2 Association classification

PathQuant can classify each association of the input by gene product into four broad categories: enzyme, transporter, other (other proteins, transcription factor, and more) and not classified using KEGG Brite database.

1.3 Metabolic network modelling

The selected KEGG pathway, encoded in KEGG XML file format (KGML), is downloaded using the KEGG API, from the most up-to-date KEGG pathways available, and then moved to a specific version folder in order to keep track, or use a specific version of interest. Users can choose which downloaded version to use or let PathQuant use the up-to-date version. The pathway is then converted into a graph of biochemical reactions (also called compound graph) with metabolites, as nodes and genes, mapped to their corresponding encoded enzymes, as edges. The topology of the pathway is captured in the constructed graph. Genes encoding enzymes catalysing multiple

reactions are mapped to multiple edges. The constructed graph represents exactly the metabolic pathways of KEGG which are built mainly with metabolites that are the main reactants of a reaction, dismissing cofactor metabolites, such as NAD, or common co-substrates/products, such as ATP or H₂O. Finally, we use a non-oriented graph as the KEGG standards are not consistent in this matter.

1.4 SRD computation

Our method computes the SRD, which is defined as the shortest reactional distance path between a given gene and a metabolite. The SRD is computed using each metabolic pathway received in input, in which a given pair is mapped. The SRD is computed as described in the workflow figure (Figure1): A distance of 0 is assigned to metabolites, which are the main substrates or products of the reaction catalysed by the enzyme encoded by the selected gene of interest. The SRDs to all other metabolites are obtained using the breadth-first search algorithm. The algorithm is used to find the SRD starting from the substrate and from the product of the mapped gene to the paired metabolite, thereby selecting the smallest SRD between these two. Figure1 depicts an example of SRD computation for a hypothetical reaction. The main reactants of this reaction: the substrate and the product are set at an SRD of 0. Everything running deeper than the substrate and the product within the graph is adding a distance of one for each depth: SRD = 1 for the green hypothetical metabolites and SRD = 2 for the blue hypothetical metabolite.

1.5 SRD metric analysis

The utility of the computed SRD metric for the annotation of gene-metabolite associations reported by mGWAS was assessed using different approaches (See Methods).

1.6 Data outputs and visualisation

PathQuant outputs a text file containing gene and metabolite classification, Enzyme Commission number (EC), KEGG Brite, KEGG IDs of used pathways for the SRD computation, and SRD values for all associations. These SRD values can also be visualised in a heatmap and global or multiple distribution plots; a few examples are available at in this manuscript (Results).

TK datasets and naming.

For the TK study, we downloaded the file named “NIHMS58114-supplement-2.xlsx” from the supplementary section of the publication. This first file contains only stringently associated pairs. We make a distinction between the TK and the “TK None” datasets as they do not use the same files from the original publication. The TK None is referring to the GWAS summary stats, without any p-value cut-off, directly downloaded from here: <https://metabolomics.helmholtz-muenchen.de/gwas/index.php?task=download>. We downloaded and merge the files named: shin_et_al.metal.out.tar.gz, shin_et_al.xeno.metal.out.tar.gz. As the positions were coming from NCBI Build 36, we performed a lift over to hg19 in order to have same build across all different studies.