



Enzyme annotation for orphan and novel reactions using knowledge of substrate reactive sites

Noushin Hadadi^{a,1}, Homa MohammadiPeyhani^{a,1}, Ljubisa Miskovic^a, Marianne Seijo^{a,2}, and Vassily Hatzimanikatis^{a,3}

^aLaboratory of Computational Systems Biology, Institute of Chemistry and Chemical Engineering, Ecole Polytechnique Fédérale de Lausanne, CH-1015 Lausanne, Switzerland

Edited by Sang Yup Lee, Korea Advanced Institute of Science and Technology, Daejeon, South Korea, and approved March 6, 2019 (received for review November 5, 2018)

Thousands of biochemical reactions with characterized activities are “orphan,” meaning they cannot be assigned to a specific enzyme, leaving gaps in metabolic pathways. Novel reactions predicted by pathway-generation tools also lack associated sequences, limiting protein engineering applications. Associating orphan and novel reactions with known biochemistry and suggesting enzymes to catalyze them is a daunting problem. We propose the method BridgIT to identify candidate genes and catalyzing proteins for these reactions. This method introduces information about the enzyme binding pocket into reaction-similarity comparisons. BridgIT assesses the similarity of two reactions, one orphan and one well-characterized nonorphan reaction, using their substrate reactive sites, their surrounding structures, and the structures of the generated products to suggest enzymes that catalyze the most-similar nonorphan reactions as candidates for also catalyzing the orphan ones. We performed two large-scale validation studies to test BridgIT predictions against experimental biochemical evidence. For the 234 orphan reactions from the Kyoto Encyclopedia of Genes and Genomes (KEGG) 2011 (a comprehensive enzymatic-reaction database) that became nonorphan in KEGG 2018, BridgIT predicted the exact or a highly related enzyme for 211 of them. Moreover, for 334 of 379 novel reactions in 2014 that were later cataloged in KEGG 2018, BridgIT predicted the exact or highly similar enzymes. BridgIT requires knowledge about only four connecting bonds around the atoms of the reactive sites to correctly annotate proteins for 93% of analyzed enzymatic reactions. Increasing to seven connecting bonds allowed for the accurate identification of a sequence for nearly all known enzymatic reactions.

reaction similarity | reactive site recognition | orphan reactions | novel (de novo) reactions | sequence similarity

Genome-scale reconstructions of metabolic networks can be used to correlate the genome with the observed physiology, though this hinges on the completeness and accuracy of the sequenced genome annotations. “Orphan” reactions, which are enzymatic reactions without protein sequences or genes associated with their functionality, are common and can be found in the genome-scale reconstructions of even well-characterized organisms, such as *Escherichia coli* (1). Recent publications reported that 40 to 50% of the enzymatic reactions cataloged in the Kyoto Encyclopedia of Genes and Genomes (KEGG) (2) lack an associated protein sequence (3, 4).

Problems with orphanlike reactions can also arise in areas such as bioremediation, synthetic biology, and drug discovery, where exploring the potential of biological organisms beyond their natural capabilities has prompted the development of tools that can generate de novo hypothetical enzymatic reactions and pathways (5–15). These de novo reactions are behind many success stories in biotechnology and can be used in the gap filling of metabolic networks (6, 12, 13, 15–18). While these enzymatic reactions have well-explained biochemistry that can conceivably occur in metabolism, they are essentially orphan reactions because they have no assigned enzyme or corresponding gene sequence. The lack of protein-encoding genes associated with the functionality of these de novo reactions limits their applicability for

metabolic engineering, synthetic biology applications, and the gap filling of genome-scale models (19). A method for associating de novo reactions to similarly occurring natural enzymatic reactions would allow for the direct experimental implementation of the discovered novel reactions or assist in designing new proteins capable of catalyzing the proposed biotransformation.

Computational methods for identifying candidate genes of orphan reactions have mostly been developed on the basis on protein sequence similarity (3, 20–22). The two predominant classes of these sequence-based methods revolve around gene/genome analysis (22–25) and metabolic information (26, 27). Several bioinformatics methods combine different aspects of these two classes, such as gene clustering, gene coexpression, phylogenetic profiles, protein interaction data, and gene proximity, for assigning genes and protein sequences to orphan reactions (28–31). All these methods use the concept of sequence similarity. Within this concept, homology between two sequences, one orphan and one well characterized, is inferred when the two share more similarity than would be expected by chance (32). Next, the biochemical function is assigned to the orphan protein sequence, assuming that homologous sequences have similar functions.

Significance

Recent advances in synthetic biochemistry have resulted in a wealth of novel hypothetical enzymatic reactions that are not matched to protein-encoding genes, deeming them “orphan.” A large number of known metabolic enzymes are also orphan, leaving important gaps in metabolic network maps. Proposing genes for the catalysis of orphan reactions is critical for applications ranging from biotechnology to medicine. In this work, the computational method BridgIT identified potential enzymes of orphan reactions and nearly all theoretically possible biochemical transformations, providing candidate genes to catalyze these reactions to the research community. The BridgIT online tool will allow researchers to fill the knowledge gaps in metabolic networks and will act as a starting point for designing novel enzymes to catalyze nonnatural transformations.

Author contributions: N.H., H.M., L.M., M.S., and V.H. designed research; N.H., H.M., and M.S. performed research; N.H., H.M., L.M., M.S., and V.H. analyzed data; and N.H., H.M., L.M., M.S., and V.H. wrote the paper.

The authors declare no conflict of interest.

This article is a PNAS Direct Submission.

This open access article is distributed under [Creative Commons Attribution-NonCommercial-NoDerivatives License 4.0 \(CC BY-NC-ND\)](https://creativecommons.org/licenses/by-nc-nd/4.0/).

Data deposition: The results of the BridgIT analysis of the KEGG 2018 orphan and novel ATLAS reactions are available at lsb-databases.epfl.ch/atlas/.

¹N.H. and H.M. contributed equally to this work.

²Present address: Institute for Environmental Sciences, Group of Environmental Physical Chemistry, Department F.-A. Forel for Environmental and Aquatic Sciences, University of Geneva, CH-1211 Geneva, Switzerland.

³To whom correspondence should be addressed. Email: vassily.hatzimanikatis@epfl.ch.

This article contains supporting information online at www.pnas.org/lookup/suppl/doi:10.1073/pnas.1818877116/-DCSupplemental.

Published online March 25, 2019.

This can be problematic because many known enzymatic activities are still missing an associated gene due to annotation errors, the incompleteness of gene sequences (33), and the fact that homology-based methods cannot annotate orphan protein sequences with no or little sequence similarity to known enzymes (3, 34). Moreover, sequence-similarity methods can provide inaccurate results because small changes in key residues could greatly alter enzyme functionality (35); in addition, it is a common observation that vastly different protein sequences can exhibit the same fold and, therefore, have similar catalytic activity, even though they look very different (36, 37).

These shortcomings motivated the development of alternative computational methods based on the structural similarity of reactants and products for identifying candidate protein sequences for orphan enzymatic reactions (31, 35, 38–42). The idea behind these approaches was to assess the similarity of two enzymatic reactions via the similarity of their reaction fingerprints; that is, the mathematical descriptors of the structural and topological properties of the participating metabolites (43), which could eliminate the problems associated with nonmatching or unassigned protein sequences. In such methods, the reaction fingerprint of an orphan reaction is compared with a set of nonorphan reference reaction fingerprints, and the genes of the most-similar reference reactions are then assigned as promising candidate genes for the orphan reaction. Reaction fingerprints can be generated based on different similarity metrics, such as the bond change, reaction center, or structural similarity (42).

One class of reaction-fingerprint computational methods compares all of the compounds participating in reactions (42), which includes both reactants and cofactors. The application of this group of methods is restricted to specific enzymatic reactions that do not involve large cofactors (31, 35, 38–42). This is because the structural information of the large cofactors overwhelmingly contributes to the corresponding reconstructed reaction fingerprint, and consequently, reactions with similar cofactors will inaccurately be classified as similar (35–38).

Another class of reaction-fingerprint methods uses the chemical structures of reactant pairs for comparison (40). While these methods can be applied to all classes of enzymatic reactions, they neglect the crucial role of cofactors in the reaction mechanism. Moreover, neither of these two classes of methods has been employed for assigning protein sequences to de novo reactions (40).

In this study, we introduce a computational method, BridgIT, that links orphan reactions and de novo reactions predicted by pathway design tools such as BNICE.ch (16), Retropath2 (15), DESHARKY (10), and SimPheny (12) with well-characterized enzymatic reactions and their associated genes. BridgIT uses reaction fingerprints to compare enzymatic reactions and is inspired by the lock-and-key principle that is used in protein docking methods (44), wherein the enzyme binding pocket is the “lock” and the ligand is a “key.” If a molecule has the same reactive sites and a similar surrounding structure as the native substrate of a given enzyme, it is then rational to expect that the enzyme will catalyze the same biotransformation on this molecule. Following this reasoning, BridgIT uses the structural similarity of the reactive sites of participating substrates together with their surrounding structure as a metric for assessing the similarity of enzymatic reactions. It is substrate-reactive-site centric, and its reaction fingerprints reflect the specificities of biochemical reaction mechanisms that arise from the type of enzymes catalyzing those reactions. BridgIT introduces an additional level of specificity into reaction fingerprints by capturing critical information about the enzyme binding pocket. More precisely, BridgIT allows us to capture approximately the 2D structure of the enzyme binding pocket by incorporating the information about sequences of atoms and bonds around the substrate reactive site.

Through several studies, we demonstrated the effectiveness of utilizing the BridgIT fingerprints for mapping novel and orphan

reactions to the known biochemistry. These reactions are mapped according to the enzyme commission (EC) (45) number, which is an existing numerical classification scheme for enzyme-based reactions. The EC number can classify enzymes at up to four levels, with a one-level classification being the most general and a four-level classification being the most specific, and these enzyme-based reactions are then represented by four numbers, one for each level, separated by periods (e.g., 1.1.1.11). We show that BridgIT is capable of correctly predicting enzymes with an identical third-level EC number, indicating a nearly identical type of enzymatic reaction, for 90% of orphan reactions from KEGG 2011 that became nonorphan in KEGG 2018. This result validates the consistency of the sequences predicted by BridgIT with the experimental observations, and it further suggests that BridgIT can provide enzyme sequences for catalyzing nearly all orphan reactions. For the remaining 10% of the orphan reactions, an in-depth sequence and structure analysis will be required to guide the sequence search and protein engineering because it is known from the enzyme analysis and classification that although reactions with common EC classification up to the third level have a nearly identical catalytic mechanism, they do not necessarily share the same sequences.

We also studied how the size of the BridgIT fingerprint impacts the BridgIT predictions. We show that BridgIT correctly identifies protein sequences using fingerprints that describe the neighborhood up to six bonds away from the atoms of the reactive site. Strikingly, we also find that it is sufficient to use the information of only three bonds around the atoms of the reactive sites of substrates to accurately identify protein sequences for 93% of the analyzed reactions.

Lastly, to indicate the utility of this computational technique, we applied BridgIT to the study of all of the 137,000 novel reactions from the ATLAS of Biochemistry, a database of all of the known and hypothetically possible biochemical reactions that connect two or more KEGG compounds (in version KEGG 2015) (46). Using our technology, we provide candidate enzymes that can potentially catalyze the biotransformation of these reactions to the research community, which should provide a basis for the engineering and development of novel enzyme-catalyzed biotransformations.

Results and Discussion

BridgIT Method. The BridgIT workflow, together with an example of its application on an orphan reaction, is demonstrated in Fig. 1. BridgIT is organized into four main steps (see *Methods* for more details): reactive site identification; reaction fingerprint construction; reaction similarity evaluation; and scoring, ranking, and gene assignment. The inputs of the workflow are (i) an orphan or a novel reaction and (ii) the collection of BNICE.ch generalized enzyme reaction rules. These reaction rules assemble biochemical knowledge distilled from the biochemical reaction databases and are used to discover de novo enzymatic reactions as well as predict all possible pathways from known compounds to target molecules (16, 46, 47). Here, we used the generalized enzyme reaction rules to extract information about the reactive sites of substrates participating in an orphan or a novel reaction, and then integrated this information into the BridgIT reaction fingerprints (Fig. 1, steps 1 and 2). We next compared the obtained BridgIT reaction fingerprints to the ones from the reference reaction database on the basis of the Tanimoto similarity scores (Fig. 1, step 3). A Tanimoto score near 0 designates reactions with no or low similarity, whereas a score near 1 designates reactions with high similarity. We used these scores to rank the assigned reactions from the reference reaction database and then identified the enzymes associated with the highest-ranked reference reactions as candidates for catalyzing the analyzed orphan or novel reaction (Fig. 1, step 4). In the following sections, we discuss the reconstructions and testing of the various components

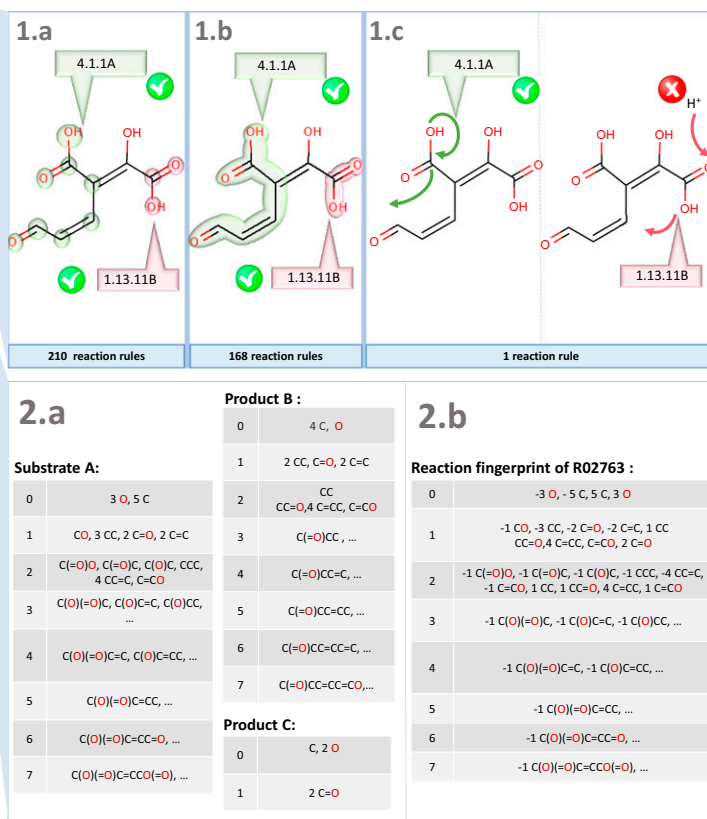
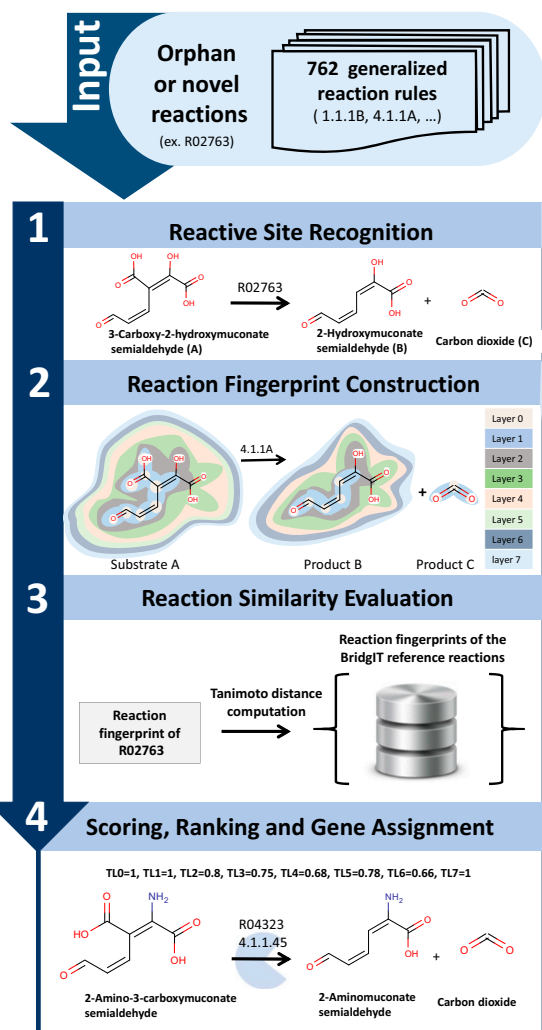


Fig. 1. Main steps of the BridgIT workflow: step 1, reactive site recognition for an input reaction (de novo or orphan); step 2: reaction fingerprint construction; step 3, reaction similarity evaluation; and step 4, sorting, ranking, and gene assignment. Steps 1.a through 1.c illustrate the procedure of the identification of reactive sites for the orphan reaction R02763. Step 1.a: two candidate reactive sites of 3-carboxy-2-hydroxyomuconate semialdehyde (substrate A) that were recognized by the rules 4.1.1. (green) and 1.13.11 (red). Step 1.b: both rules recognized the connectivity of atoms within two candidate reactive sites. Step 1.c, only reaction rule 4.1.1. can explain the transformation of substrate A to products. Step 2.a shows the fragmentation of reaction compounds, whereas step 2.b illustrates the mathematical representations of the corresponding BridgIT reaction fingerprints.

of BridgIT as well as the results of our main analyses. A web tool of BridgIT can be consulted at lcsb-databases.epfl.ch/pathways/BridgIT.

Reference Reaction Database. The BridgIT reference reaction database is an essential component of the BridgIT workflow (Fig. 1). It consists of well-characterized reactions with associated genes and protein sequences and was built based on the KEGG 2016 reaction database (see *Methods*). The KEGG database is the most comprehensive database of enzymatic reactions and provides information about biochemical reactions together with their corresponding enzymes and genes. However, half of KEGG reactions lack associated genes and protein sequences and are thus considered to be orphan reactions. The BridgIT reference database was built using the KEGG reactions that (i) can be reconstructed by the existing BNICE.ch generalized reaction rules and are elementally balanced (5,270 reactions) and (ii) are nonorphan (5,049 reactions). This restriction removes reactions that lack characterized substrate reactive sites, meaning that they cannot be used in our comparisons. As a result, the reference reaction database contains information for 5,049 of 9,556 KEGG reactions (Dataset S1, Table S1).

Sensitivity Analysis of the BridgIT Fingerprint Size. The defining characteristic of the BridgIT reaction fingerprint is that it is centered around the reactive site of the reaction substrate(s). The number of description layers in the BridgIT fingerprint—the fingerprint size—defines how large a chemical structure around the reactive site we consider when evaluating the similarity (see *Methods*). To investigate to what extent the fingerprint size affects the similarity results, we performed a sensitivity analysis in which we varied the fingerprint size between 0 and 10.

For this analysis, we considered the 5,049 nonorphan KEGG reactions that existed in the BridgIT reference reaction database. We started by forming reaction fingerprints that contained only the description layer 0 (fingerprint size 0) and evaluated how many of 5,049 nonorphan reactions BridgIT could correctly identify. That is, we evaluated whether the BridgIT algorithm with these reaction fingerprints could map each of these reactions to itself. We next formed the reaction fingerprints using only the description layers 0 and 1 (fingerprint size 1), and we performed the evaluation again. We repeated this procedure until the final step, in which we formed the reaction fingerprints with 10 description layers (fingerprint size 10).

As expected, the increase in the fingerprint size (i.e., specificity) led to a decrease in the average number of similar reactions assigned to the studied reactions. Moreover, the more description layers that were incorporated into the BridgIT fingerprint, the more accurately BridgIT matched the analyzed reactions (Table 1). For the fingerprint size 7, BridgIT correctly mapped 100% of the analyzed reactions; that is, each of the 5,049 nonorphan reactions was matched to itself in the reference reaction database. This indicated that the information about chains of eight atoms along with their connecting bonds around the reactive sites was sufficient for BridgIT to correctly match all nonorphan KEGG reactions, and we chose the fingerprint size 7 for our further studies.

BridgIT Reaction Fingerprints Offer Improved Predictions. To evaluate BridgIT performances against existing approaches in this field (40, 42, 48), we performed two comparative studies. In the first study, we repeated the analysis from the previous section using the standard reaction difference fingerprint (*Methods* and *SI Appendix, Fig. S4*), which is used and discussed in detail in structure similarity methods such as RxnSim (38) and RxnFinder (39), to assess the benefits of introducing the information about the reactive site of substrates into the reaction fingerprints. A comparison of the two sets of predictions on 5,049 nonorphan reactions showed that the predictions obtained with BridgIT-modified fingerprints were significantly better than those obtained with the standard ones. BridgIT identified 100% of nonorphan reactions correctly versus the 71% success rate for the standard fingerprint method (*Dataset S1, Table S4*). Furthermore, BridgIT correctly matched 93% of the analyzed enzymatic reactions using the information about only four connecting bonds around the atoms of the reactive sites (fingerprint size 4) (Table 1), which exceeds the 71% of matched reactions when using the standard reaction fingerprints (fingerprint size 7).

The inferior performance of the standard reaction fingerprint method arose from three main sources. First, fragments from the substrate and product sets were cancelled out upon algebraic summation inside the fingerprint description layers (see *Methods*), in which description layers 0 and 1 define the single atoms and the connected pairs of atoms of the reactive site, and layers 2 to 7 include information about the chemical structure around the reactive site that contains up to eight atoms and seven bonds (Fig. 1). This cancellation occurred in all description layers (fingerprint size 7) for 246 nonorphan reactions—that is, their standard fingerprints were empty (*Dataset S1, Table S3*). As an example, Fig. 2A shows the standard reaction fingerprint of KEGG reaction R00722 that was empty for the standard fingerprint method. The information about reactive sites introduced in the BridgIT reaction fingerprints prevents such cancellations, since BridgIT does not include the atoms of the reactive site(s) in the process of the algebraic summation of the substrate and product

set fragments (see *Methods*). As a result, BridgIT mapped R00722 to itself and identified R00330 as the most similar reaction to R00722 (Fig. 2A). Indeed, according to the KEGG database, the enzyme 2.7.4.6 catalyzes both reactions.

Second, the performance of the standard reaction fingerprint suffered because the first description layer of the standard fingerprint was empty for an additional 1,129 reactions, which indicated that these fingerprints did not represent the bond changes during the reaction (*Dataset S1, Table S4*).

Third, the remaining 89 mismatched nonorphan reactions had partial cancellations in the fingerprint description layers. For example, the standard fingerprint method incorrectly identified R03132 as the most similar to R00691, whereas BridgIT identified R00691 and R01373 as the most similar to R00691 (Fig. 2B), which matches the KEGG reports indicating that both R00691 and R01373 can be catalyzed by either EC 4.2.1.51 or EC 4.2.1.91.

In the second study, we compared the performance of BridgIT method against three state-of-the-art methods—EC-BLAST (42), Selenzyme (48), and E-zyme2 (40)—on three benchmark problems. The first two benchmark problems consisted of identifying the most-similar reactions to two example reactions, each representing a class of reactions that appear ubiquitously in biochemical networks. We chose R00722 (Fig. 2) to exemplify the first class of reactions characterized by a very similar structure of substrates and products, and chose R07500 to represent the class of multisubstrate multiproduct reactions (*SI Appendix, Tables S1 and S2*). The third benchmark problem represented the intermolecular transferases (EC 5.4.4) that catalyze the transfer of a hydroxyl group to another part of a molecule. Similar to other isomerases, in this class of reactions, the substrate and the product have the same chemical formula but different bond connectivity. We chose R09708 to exemplify this class of reactions (*SI Appendix, Table S3*).

For the three benchmark reactions, we ranked the similar reactions proposed by each of the methods according to the corresponding similarity scores, and the top 100 similar reactions proposed by each method were used for comparisons.

The most similar reaction proposed by BridgIT correctly matched the fourth-level EC number (2.7.4.6) of the first benchmark reaction R00722 (*SI Appendix, Table S1*). Three of four EC-BLAST variants (42) proposed a set of the reactions with the maximal similarity score (*SI Appendix, Table S1*). This set contained not only reactions that correctly matched the fourth-level EC number of R00722, but also reactions with EC numbers not even matching the first-level EC number of the benchmark reaction (*SI Appendix, Table S1*). The three variants of Selenzyme (48) proposed reactions that could match only the third-level EC number of R00722, whereas E-zyme2 (40) was unable to find a matching reaction due to very similar structures in the substrate–product pairs (*SI Appendix, Table S1*).

In the second benchmark, none of the investigated methods could propose reactions that match the EC number of R07500 (2.5.1.115) up to the fourth level, and all methods could match the third-level EC number for this reaction (*SI Appendix, Table S2*). BridgIT proposed 39 similar reactions matching the third-level EC numbers of R07500, whereas the EC-BLAST variant with structural similarity proposed 45 similar reactions, Selenzyme proposed 10, E-zyme2 proposed 9, and the three other EC-BLAST variants proposed 5 to 7 (*SI Appendix, Table S2*). In addition, we performed receiver operating characteristic (ROC) analysis on the sets of proposed similar reactions, and of all the compared methods, BridgIT had the highest area under the ROC curve (AUC) index of 0.95, meaning that it had the best performance among the compared methods for this class of reactions (*SI Appendix, Table S2*).

In the third benchmark, BridgIT was the only method that could match the third-level EC number of R09708 (5.4.4.4). It proposed linalool isomerase (5.4.4.8) to catalyze this reaction, and remarkably, it was reported in the literature that this enzyme could catalyze stereospecific isomerization of (3*S*)-linalool to

Table 1. Correctly mapped reactions as a function of the size of the BridgIT fingerprint

Fingerprint size	Correctly mapped reactions, %
0	4.3
1	35.2
2	60.5
3	72.1
4	92.7
5	97.8
6	98.6
7	100
8	100
9	100
10	100

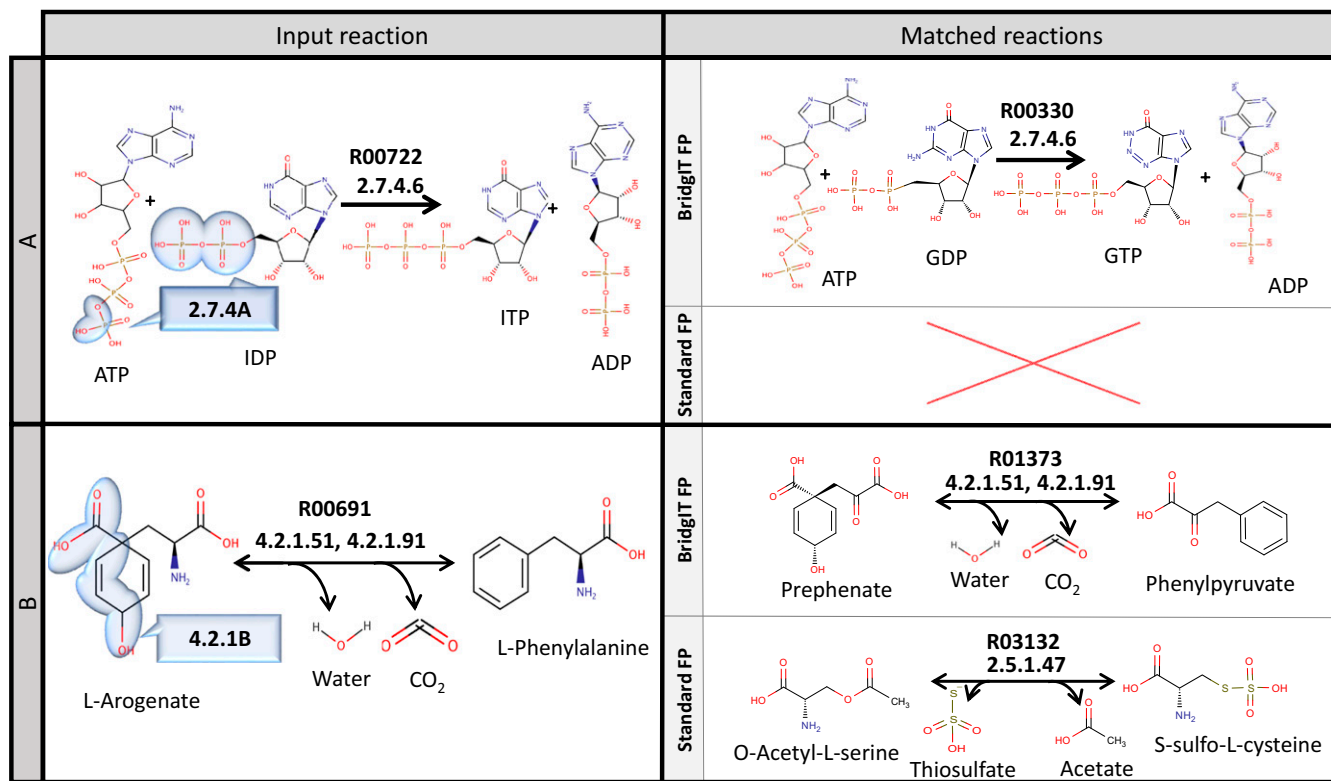


Fig. 2. Comparison of the results obtained with the BridgIT and standard fingerprint on two example KEGG reactions. (A) The input reaction R00722 (Left) and the most-similar reactions (Right) identified with the BridgIT and standard fingerprints. Note that the standard fingerprinting method failed to find a similar reaction to R00722 due to cancellations inside all fingerprint description layers. (B) The input reaction R00691 (Left) and the most-similar reactions (Right) identified with the BridgIT and standard fingerprints.

geraniol (49). Other methods proposed catalyzing enzymes from EC class 4 (lyases) (*SI Appendix, Table S3*). Moreover, only BridgIT and the EC-BLAST variants with structural similarity and bond change could capture the structural changes in R09708, whereas E-zyme2, Selenzyme, and the two remaining variants of EC-BLAST could not map this reaction to itself (*SI Appendix, Table S3*).

The results of these three studies demonstrate the potential of BridgIT to outperform the currently available methods for enzyme annotation.

From Reaction Chemistry to Detailed Enzyme Mechanisms. Approximately 15% of KEGG reactions (1,532 reactions) are assigned to more than one enzyme and EC number; that is, multiple enzymes can catalyze a specific biotransformation through different enzymatic mechanisms. For example, KEGG reaction R00217 is assigned to three different EC numbers, 4.1.1.3 (oxaloacetate carboxy-lyase) and 1.1.1.38 (both malate dehydrogenases), and the corresponding reactions involve different mechanisms (Fig. 3). The reaction mechanism of the 4.1.1.3 enzyme is well understood, as it belongs to the carboxy-lyases in which a carbon-carbon bond is broken and a molecule of CO₂ is released. This enzyme can decarboxylate three different compounds: glutaconyl-CoA, methylmalonyl-CoA, and oxaloacetate (from this example). The overlapping reactive site of these three compounds is captured in the 4.1.1B rule of BNICE.ch (Fig. 3C). In contrast, the 1.1.1.38 enzyme found in bacteria and insects and the 1.1.1.40 enzyme found in fungi, animals, and plants are rather specific enzymes that decarboxylate oxaloacetate and malate with two different mechanisms. The decarboxylation is performed with (in the case of malate) or without (in the case of oxaloacetate) the incorporation of NAD⁺ as a cofactor. The only difference in the structure of these two molecules is in having ei-

ther a ketone or an alcohol group on the second carbon. Consequently, the structure of the reactive site that these enzymes recognize has to reflect the difference between malate and oxaloacetate, and this is well captured in the 1.1.1A rule of BNICE.ch. The 4.1.1B rule requires a less specific reactive site compared with the 1.1.1A rule, and these two rules have two different reaction




Fig. 3. A multienzyme reaction such as R00217 can be catalyzed by more than one enzyme. BridgIT identified two distinct fingerprints for this reaction that correspond to two reactive sites of oxaloacetate (A). The reactive site recognized by the 1.1.1 rule (B) is more specific (blue substructure) than the one recognized by the 4.1.1 rule (C) (green substructure).

fingerprints for catalyzing the same reaction R00217 because they describe different mechanisms for the same reaction.

Moreover, for 42% of the KEGG reactions that have a single enzyme assigned to them, BNICE.ch identified multiple alternative reactive sites and created multiple reaction fingerprints that describe the biotransformation of these reactions. Therefore, a single reaction from KEGG was translated into more than one fingerprint in the BridgIT reference database. This way, by preserving the information about enzyme binding pockets, the reconstructed BridgIT reference reaction database expands from 5,049 reactions to 17,657 reaction fingerprints corresponding to 17,657 detailed reaction mechanisms.

Currently, BridgIT is the only method that can distinguish different reaction mechanisms for the reactions catalyzed by different enzymes. As a consequence, BridgIT can propose distinct sets of protein sequences corresponding to distinct mechanisms and rank them according to the BridgIT score. The protein sequences can then be prioritized based on the BridgIT ranking, enzyme specificity, and the host organism.

Comparison of BridgIT and BLAST Predictions. As a means to relate reaction structural similarity obtained using BridgIT with reaction sequence similarity obtained using BLAST (50), we applied these two techniques in parallel on a subset of reactions and their corresponding protein sequences from the reference reaction database. We compared the similarity results of BridgIT with those of BLAST and statistically assessed BridgIT performance using ROC curve analysis (*SI Appendix, Figs. S1 and S2*).

We chose *E. coli* BW29521 (EBW) as our benchmark organism for this analysis. There were 531 nonorphan reactions in EBW associated with 413 protein sequences. In total, there were 731 reaction–gene associations (*Dataset S1, Table S2*), as there were reactions with more than one associated gene and genes associated with more than one reaction. We removed all the nonorphan reactions of EBW from the BridgIT reference database and removed their associated protein sequences from the KEGG protein sequence database (*Dataset S1, Table S2*). We then used BridgIT to assess the structural similarity of the 531 EBW reactions to the BridgIT reference reactions using the Tanimoto score, and we applied BLAST to quantify the similarity of the 413 EBW protein sequences to the protein sequences of reactions from the BridgIT reference database using e-values. The concept of the validation procedure is illustrated in *SI Appendix, Fig. S1*. We provide a list of BridgIT reaction–reaction comparisons together with BLAST sequence–sequence comparisons (*Dataset S1, Table S2*).

Comparing Reaction (BridgIT) and Sequence (BLAST) Similarity Scores. We considered two sequences to be similar if BLAST reported an e-value of less than 10^{-10} for their alignment. For a chosen discrimination threshold (DT) of the global Tanimoto score (T_G), we considered the BridgIT prediction of similarity between an EBW reaction and a BridgIT reference reaction with a T_G score as (i) true positive (TP) if $T_G > DT$ and their associated sequence (s) were similar (e-value $< 10^{-10}$); (ii) true negative (TN) if not similar for both BridgIT ($T_G < DT$) and BLAST+ (e-value $> 10^{-10}$); (iii) false positive (FP) if similar for BridgIT ($T_G > DT$) but not similar for BLAST+ (e-value $> 10^{-10}$); and (iv) false negative (FN) if not similar for BridgIT ($T_G < DT$) but similar for BLAST+ (e-value $< 10^{-10}$).

We then counted the number of TPs, TNs, FPs, and FNs for all 531 reactions and summed these quantities to obtain the total number of TPs, TNs, FPs, and FNs per chosen DT. We repeated this procedure for a set of DT values varying across the interval between 0 and 1. Lastly, we used the total number of TPs, TNs, FPs, and FNs to compute the TP and FP rates for the ROC curve analysis (*SI Appendix, Fig. S2A*). The ROC curve indicated that the reaction comparison based on reaction structural similarity

(BridgIT) was comparable to the one based on reaction sequence similarity (BLAST). Indeed, the obtained AUC score for the BridgIT classifier was 0.91, indicating that the similarities between the two methods were very high (*SI Appendix, Fig. S2A*). We next studied whether the type of compared reactions affected the accuracy of BridgIT predictions by categorizing reactions according to their first-level EC class, which indicates the broadest category of enzyme functionality, and then performing the ROC analysis for each class separately (*SI Appendix, Fig. S2A*). The analysis revealed that BridgIT performed well with all major enzyme classes, as represented by the high AUC scores ranging from 0.88 (EC 1) to 0.96 (EC 5).

We next analyzed the accuracy of BridgIT classification as a function of the DT of the Tanimoto score (*SI Appendix, Fig. S2B*). The accuracy ranged from 43% for a DT value of 0.01 to 85% for a DT value of 0.30. For DT values > 0.30 , the accuracy monotonically decreased toward 62% for a DT value of 1. The classifier was overly conservative for DT values > 0.30 and was rejecting TPs (*SI Appendix, Fig. S2B*). More specifically, for a DT value of 0.30, the TP percentage was 38%, whereas for a DT value of 1, it was reduced to 3%. In contrast, the TN percentage increased very slightly for DT values > 0.30 , whereas for a DT value of 0.30, it was 46%, and for a DT value of 1, it was 57% (*SI Appendix, Fig. S2B*). Based on this analysis, we chose 0.30 as an optimal DT value for further studies.

A sensitivity analysis of BridgIT results to the variations in the e-value threshold ranging from 10^{-10} to 10^{-50} is provided in *SI Appendix, Fig. S3*.

BridgIT Analysis of Known Reactions with Common Enzymes. The 5,049 reactions in the reference database were catalyzed by only 2,983 enzymes; that is, there were promiscuous enzymes that catalyzed more than one reaction. Of the 2,983 enzymes, 844 were promiscuous, catalyzing 2,432 of the reactions (*Dataset S1, Table S5*). Interestingly, BridgIT correctly assigned more than 80% of these 2,432 reactions to their corresponding promiscuous enzyme. An example of such a group is given in Table 2. This table shows the same enzymes listed across the top and down the left side, with the corresponding Tanimoto scores indicating the accuracy of BridgIT's classifications. The overall high scores in this table indicate the accuracy of the enzyme assignments.

We investigated the remaining 20% of reactions in depth, and we observed that the Tanimoto scores of the first two description layers (see *Methods*) indicated a very low similarity between the reactions catalyzed by the same enzyme. This result suggested that such enzymes were either multifunctional (i.e., they had more than one reactive site) (Fig. 4) or were incorrectly classified in the EC classification system.

BridgIT Validation Against Biochemical Assays. To assess BridgIT's performance using biochemically confirmed reactions, we performed two validation studies on sets of orphan (study I) and novel (study II) reactions. Since the known reactions in KEGG are all experimentally confirmed using biochemical assays, we

Table 2. A group of five reactions catalyzed by enzyme 1.1.1.219

Catalyzed reactions	R03123	R03636	R05038	R07999	R07998
R03123	1	0.96	0.93	0.93	0.98
R03636	0.96	1	0.96	0.94	0.95
R05038	0.93	0.96	1	0.97	0.91
R07999	0.93	0.94	0.97	1	0.91
R07998	0.98	0.95	0.91	0.91	1

Values are Tanimoto scores given for the comparison between the reaction listed across the top and the reaction listed down the left side of the table.

A multi-functional enzyme: EC 1.11.1.8




Fig. 4. Multifunctional enzymes can catalyze reactions with two different reactive sites. R03539 (A) and R03208 (B) are catalyzed by the same enzyme, 1.11.1.8. However, the reactive sites of these substrates are completely different.

could use this pooled experimental data from hundreds of laboratories to demonstrate BridgIT's ability to identify potential enzymes for catalyzing the biologically relevant orphan reactions on a large scale.

Study I. We compared the number of orphan reactions in two versions of the KEGG reaction database: KEGG 2011 and KEGG 2018. We found that 234 orphan reactions from KEGG 2011 were later associated with enzymes in KEGG 2018, meaning that they became nonorphan reactions (Dataset S1, Tables S6–S8). Since these newly classified reactions have been experimentally confirmed, we used these 234 reactions as a benchmark to evaluate BridgIT's performance.

We formed the reference reaction database using the reactions from KEGG 2011 (see Methods), and we compared the BridgIT results with the KEGG 2018 enzyme assignments up to the third EC level. Remarkably, BridgIT- and KEGG 2018-assigned enzymes matched to the third EC level for 211 of 234 (90%) reactions (Dataset S1, Tables S6 and S7). This means that BridgIT accurately predicted the enzyme mechanism for enzymes that have been biochemically confirmed to catalyze a large majority of the orphan reactions in 2011. In addition, the set of protein sequences proposed by BridgIT comprised highly related protein sequences to the ones assigned to these enzymes in KEGG 2018.

The 234 reactions are catalyzed by 168 enzymes with specified fourth-level EC numbers in KEGG 2018. However, only 29 of these 168 enzymes were cataloged in KEGG 2011, and the remaining 139 enzymes had new fourth-level EC classes assigned in KEGG 2018, meaning that BridgIT had access to only the 29 enzymes that were classified in KEGG 2011 from which the reference reaction database was built. The 29 enzymes catalyzed 35 of the 234 studied reactions. For 29 of these 35 (83%) orphan reactions, the BridgIT algorithm predicted the same sequences that KEGG 2018 assigned to these reactions (Dataset S1, Table S9). A higher matching score when comparing up to the third EC level rather than the fourth EC level is likely because BridgIT uses BNICE.ch generalized reaction rules, which describe the biotransformations of reactions with specificities up to the third EC level.

Study II. The ATLAS of Biochemistry (46) provides a comprehensive catalog of theoretically possible biotransformations between KEGG compounds and can be mined for novel biosynthetic routes for a wide range of applications in metabolic engineering,

synthetic biology, drug target identification, and bioremediation (40). We studied the 379 reactions from the ATLAS of Biochemistry that were novel in KEGG 2014 and were later experimentally identified and cataloged in KEGG 2018.

We formed the reference reaction database using the reactions from KEGG 2014 and applied BridgIT to these 379 reactions. For 334 of these 379 reactions, BridgIT proposed similar known reactions with a Tanimoto score higher than 0.30, thus providing promising protein sequences for enzymes catalyzing these reactions (Dataset S1, Table S10). For 14 of these novel reactions, BridgIT assigned the same sequences that were assigned in KEGG 2018 (Dataset S1, Table S11). An example of such a reaction is rat132341, which was a novel reaction in 2014 and later cataloged as R10392 in KEGG 2018 (Fig. 5A). The BridgIT analysis of this reaction revealed that R03444, which is catalyzed by enzyme 4.2.1.114, is the structurally closest reaction to this novel one, suggesting that protein sequences from EC 4.2.1.114 can catalyze this novel reaction. This was later confirmed by experimental biochemical evidence, as R10392 is associated with the same EC 4.2.1.114 enzyme in KEGG 2018. There are 243 available protein sequences for enzyme 4.2.1.114, and one sequence already has a confirmed protein structure (Fig. 5C). Therefore, BridgIT results were validated using experimental biochemical evidence on a large scale.

BridgIT Predictions for KEGG 2018 Orphan Reactions. We applied BridgIT to the 810 orphan KEGG 2018 reactions that could be reconstructed using the BNICE.ch generalized reaction rules. The remaining 1,646 orphan reactions could not be reconstructed because they are not balanced or they lack the structure for at least one of their substrates. Remarkably, BridgIT identified corresponding reference reactions with Tanimoto scores higher than the optimal threshold value of 0.30 for 97% of the orphan reactions. The remaining 3% of orphan reactions had a low similarity with the reference reactions. A large number of the orphan reactions originate from the pathways toward plant and microbial natural products that frequently involve complex and less-investigated classes of enzymes such as polyketide synthases (PKSs), nonribosomal peptide synthetases (NRPSs), terpene cyclases (TCs), and cytochromes P450 (CYPs). Interestingly, BridgIT mapped 112 of 810 orphan reactions back to these families: It predicted that 72 orphan reactions can be catalyzed by CYPs, 33 by PKSs, 6 by NRPSs, and 1 by TC (Dataset S1, Tables S12–S15).

This result and the fact that BridgIT correctly mapped 100% of nonorphan KEGG reactions suggests that as our knowledge of biochemistry expands, the annotation of novel and orphan reactions using tools such as BridgIT will also improve.

BridgIT Predictions for ATLAS Novel Reactions. We further utilized BridgIT to identify candidate enzymes for all the 137,000 de novo, orphanlike ATLAS reactions. These candidate enzymes can be used directly in systems biology designs if the matched enzymes perform the desired catalysis or if their amino acid sequences can be optimized through protein engineering to achieve the desired results. We found that 7% of novel ATLAS reactions were matched to known KEGG reactions with a Tanimoto score of 1 (perfect match), while 88% were similar to KEGG reactions with a Tanimoto score higher than the optimal threshold value of 0.30. Therefore, BridgIT could identify promising enzyme sequences for catalyzing 95% of novel ATLAS reactions. The remaining 5% of these reactions were not similar to any of the well-characterized known enzymatic reactions.

Finding well-characterized reactions that are similar to novel ones is crucial for evolutionary protein engineering as well as computational protein design, and methods like BridgIT can be instrumental in moving from a concept to the experimental implementation of de novo reactions. Additionally, to facilitate

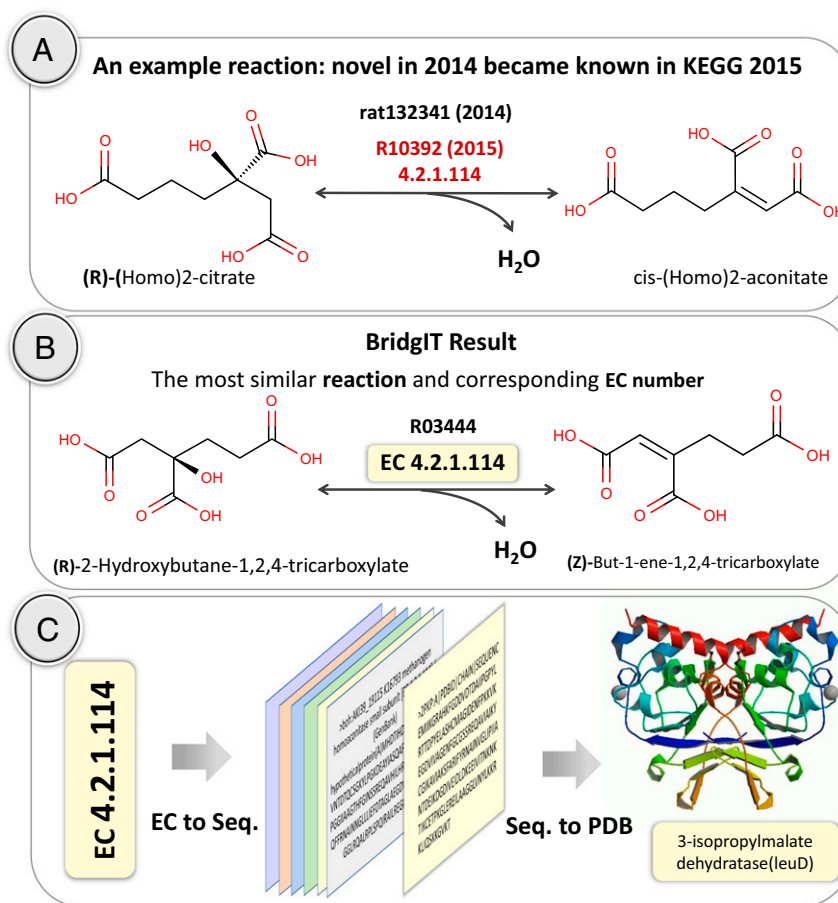


Fig. 5. Details of the BridgIT verification procedure that was performed on ATLAS reaction rat132341, which was novel in KEGG 2014 and later experimentally identified and cataloged in KEGG 2018—that is, it became a nonorphan reaction (R10392). (A) rat132341 catalyzes the conversion of (R)-(homo)2-citrate to *cis*-(homo)2-aconitate. (B) Using the biochemical knowledge of KEGG 2014, BridgIT predicts the KEGG reaction R03444, which is catalyzed by a 4.2.1.114-class enzyme, as the most similar known reaction to rat132341. Remarkably, the same enzyme is later assigned to R10392 in KEGG 2018 with the corresponding biochemical confirmation. (C) The identified EC number (4.2.1.114) can be used to extract the corresponding protein sequences along with their crystal structures.

the experimental implementation of novel ATLAS reactions in metabolic engineering, in systems and synthetic biology, and in bioremediation studies, we can use the BridgIT similarity scores as confidence measures for evaluating the feasibility.

The results of the BridgIT analysis of the KEGG 2018 orphan and novel ATLAS reactions are available on the website lcsb-databases.epfl.ch/atlas/.

Methods

In BridgIT, the Tanimoto score is used to quantify the similarity of reaction fingerprints. BridgIT allows us to do the following: (i) compare a given novel or orphan reaction to a set of reactions that have associated sequences, subsequently referred to as the reference reactions; (ii) rank the identified similar reactions based on the computed Tanimoto scores; and (iii) propose the sequences of the highest ranked reference reactions as possible candidates for encoding the enzyme of the given *de novo* or orphan reaction.

Reactive Site Identification. An enzymatic reaction occurs when its substrate(s) fits into the binding site of an enzyme. Since the structure and geometry of the binding sites of enzymes are complex and most of the time not fully characterized, we proposed focusing on the similarity of the reactive sites of their substrates. Following this, we used the expert-curated, generalized reaction rules of BNICE.ch to identify the reactive sites of substrates. These reaction rules have third-level EC identifiers (e.g., EC 1.1.1) and encompass the following biochemical knowledge of enzymatic reactions: (i) information about atoms of the substrate's reactive site; (ii) information about connectivity (atom-bond-atom); and (iii) exact information of bond breakage and formation during

the reaction. As of July 2017, BNICE.ch contains 381 bidirectional generalized reaction rules that can reconstruct 6,528 KEGG reactions (46).

Given a novel or orphan reaction, the reactive sites of its substrate(s) are identified in three steps. In the first step, the BNICE.ch generalized reaction rules that can be applied to groups of atoms from the analyzed substrates are identified, and then the information about the identified rules and the corresponding groups of atoms is stored. Subsequently, these groups of atoms are referred to as the candidate substrate reactive sites. In the second step, among the identified rules, only the ones that can recognize the connectivity between the atoms of the candidate substrate reactive sites are kept. In the third step, whether the biotransformation of a substrate(s) to a product(s) can be explained by the rules retained after the second step is tested. The candidate reactive sites corresponding to the rules that have passed the three-step test are validated and used for the construction of reaction fingerprints.

We illustrate this procedure on an orphan reaction R02763, which catalyzes the conversion of 3-carboxy-2-hydroxymuconate semialdehyde (substrate A) to 2-hydroxymuconate semialdehyde and carbon dioxide (Fig. 1). In the first step, 210 rules were identified out of 361 rules that could be applied to groups of atoms of substrate A (Fig. 1, step 1.a). Of the 210 rules, 168 matched the connectivity (Fig. 1, step 1.b). Lastly, the 168 reaction rules were applied to substrate A for bond breaking and formation comparisons, and one rule could explain the transformation of substrate A to the products (Fig. 1, step 1.c).

Reaction Fingerprint Construction. Molecular fingerprints, which are the linear representations of the structures of molecules, have been used in many methods and for different applications, especially for structural comparison of compounds (51, 52). One of the most commonly used molecular fingerprints is the Daylight fingerprint (51), which decomposes a molecule into

eight layers starting from layer zero that accounts only for atoms. Layer 1 expands one bond away from all the atoms and accounts for atom-bond-atom connections. This procedure is continued until layer 7, which includes seven connected bonds from each atom. There are two types of Daylight reaction fingerprints: (i) structural reaction fingerprints, which are simple combinations of reactant and product fingerprints, and (ii) reaction difference fingerprints, which are the algebraic summation of reactant and product fingerprints multiplied by their stoichiometry coefficients in the reaction. In this study, we propose a modified version of the reaction difference fingerprint. The procedure for formulating BridgIT reaction fingerprints is demonstrated through an example reaction (Fig. 1, step 2).

Starting from the atoms of the identified substrate reactive site, eight description layers of the molecule were formed, where different layers consisted of fragments with different lengths. Fragments were composed of atoms connected through unbranched sequences of bonds. Depending on the number of bonds included in the fragments, different description layers of a molecule were formed as follows:

Layer 0: Describes the type of each atom of the reactive site together with its count. For example, the substrate of the example reaction at layer 0 was described as three oxygens and five carbon atoms (Fig. 1, step 2.a).

Layer 1: Describes the type and count of each bond between pairs of atoms in the reactive site. In the example, the substrate at layer 1 was described with six fragments of length 1: one C–O, three C–C, two C=O, and one C=C bond (Fig. 1, step 2.a). Fragments are shown by their simplified molecular-input line-entry system (SMILES) molecular representation (53). To convert SMILES to canonical SMILES, we used Open Babel C++ library (52).

Layer 2: Describes the type and count of fragments with three connected atoms. While layers 0 and 1 described the atoms of reactive sites, starting from layer 2, atoms that were outside the reactive site were also described. In the illustrated example, there were six different fragments of this type (Fig. 1, step 2.a).

The same procedure was used to describe the molecules up to layer 7. Interestingly, and consistent with the previously reported result (43), we found that the seven-layer description was good enough to capture the structure of most of the metabolites in biochemical reactions, therefore providing a precise reaction fingerprint. Note that not all description layers are needed to describe less complex molecules. For example, product C (carbon dioxide) was fully described using only layer 0 and layer 1 (Fig. 1, step 2.a). For very large molecules, the description layers that contain fragments with more than eight connected atoms can be used.

For each layer, the substrate set was formed by merging all the fragments and their type and count in the substrate molecules of the reaction, and the product set was formed by merging all the fragments (type and count) in the product molecules of the reaction. In both sets, the count of each fragment was multiplied by the stoichiometric coefficients of the corresponding compound in the reaction. Lastly, the reaction fingerprints were created by summing the fragments of the substrate and product sets for each layer (Fig. 1, step 2.b).

Introducing the specificity of reactive sites into the reaction fingerprint allows BridgIT to capitalize on the information about enzyme binding pockets (16). To keep this valuable information throughout the generation of reaction fingerprints, BridgIT does not consider the atoms of the reactive site(s) when performing the algebraic summation of the substrate and product set fragments. Consequently, the BridgIT algorithm enables retaining, tracking, and emphasizing the information of the reactive site(s) in all the layers of the reaction fingerprint, which distinguishes it from the existing methods.

Reaction Similarity Evaluation. The similarity of two reactions was quantified using the similarity score between their fingerprints, subsequently referred to as reaction fingerprints A and B. In this study, the Tanimoto score, which is an extended version of the Jaccard coefficient and cosine similarity, was used (54). Values of the Tanimoto scores near 0 indicate reactions with no or negligible similarity, whereas values near 1 indicate reactions with high similarity.

The Tanimoto score for each descriptive layer, T_{Lk} , together with the global Tanimoto score, T_G , was calculated. The Tanimoto score for the k -th descriptive layer was defined as

$$T_{Lk} = \frac{c_k}{a_k + b_k - c_k}$$

where a_k was the count of the fragments in the k -th layer of reaction fingerprint A; b_k was the count of the fragments in the k -th layer of reaction fingerprint B; and c_k was the number of common k -th layer fragments of

reaction fingerprints A and B. Two fragments were equal if their canonical SMILES and their stoichiometric coefficients were identical. The global Tanimoto similarity score, T_G , was defined as follows:

$$T_G = \frac{\sum_{k=0}^7 c_k}{\sum_{k=0}^7 a_k + \sum_{k=0}^7 b_k - \sum_{k=0}^7 c_k}$$

For each reaction fingerprint, its Tanimoto similarity score was calculated against the reaction fingerprints from the BridgIT reference database, which contained reaction fingerprints of all known, well-characterized enzymatic reactions (Fig. 1, step 3).

Sorting, Ranking, and Gene Assignment. For a given input reaction, the reference reactions were ranked using the computed T_G scores. The algorithm distinguished between the identified reference reactions with the same T_G score based on the T_L score of layers 0 and 1 and allowed the user to assign ranking weights to specified layers. The protein sequences associated with the highest ranked (i.e., the most similar) reference reactions were then assigned to the input reaction (Fig. 1, step 4).

Conclusions

We developed the computational tool BridgIT to evaluate and quantify the structural similarity of biochemical reactions by exploiting the biochemical knowledge of BNICE.ch generalized reaction rules. Because the generalized reaction rules can identify reactive sites of substrates, BridgIT can translate the structural definition of biochemical reactions into a type of reaction fingerprint that explicitly describes the atoms of the substrates' reactive sites and their surrounding structure. Through the analysis of 5,049 known and well-defined biochemical reactions, we found that knowledge of the neighborhood up to three bonds away from the atoms of the reactive site can predict biochemistry and match catalytic protein sequences. The reaction fingerprints proposed in this work can be used to compare all novel and orphan reactions to well-characterized reference reactions and, consequently, to link them with genes, genomes, and organisms. We demonstrated through several examples the improvements that the BridgIT fingerprint brings to the field compared with the fingerprints currently existing in the literature.

A drawback of traditional sequence-similarity methods is that they cannot identify protein sequence candidates for de novo reactions, which we have shown BridgIT can do.

We tested BridgIT predictions against experimental biochemical evidence, within two large-scale validations studies on sets of 234 orphan and 379 de novo reactions. The reactions from these two sets were unknown in the previous versions of the KEGG database but were later experimentally confirmed and cataloged in KEGG 2018. BridgIT predicted the exact or a highly related enzyme for 89% of these reactions.

We further applied BridgIT to the entire catalog of de novo reactions of the ATLAS of Biochemistry database and proposed several candidate enzymes for each of them. The candidate enzymes for these de novo reactions can either be immediately capable of catalyzing these reactions or serve as initial sequences for enzyme engineering. The obtained BridgIT similarity scores can also be used as a confidence score to assess the feasibility of the implementation of novel ATLAS reactions in metabolic engineering and systems biology studies.

The applications of BridgIT go beyond merely bridging gaps in metabolic reconstructions: This method can be used to identify the potential utility of existing enzymes for bioremediation as well as for various applications in synthetic biology and metabolic engineering. As the field of metabolic engineering grows and metabolic engineering applications increasingly turn toward the production of valuable industrial chemicals such as 1,4-butanediol (55, 56), we expect that methods for the design of de novo synthetic pathways, such as BNICE.ch (16), and methods for identifying candidate enzymes for de novo reactions, such as BridgIT, will grow in importance.

1. Orth JD, et al. (2011) A comprehensive genome-scale reconstruction of Escherichia coli metabolism—2011. *Mol Syst Biol* 7:535.
2. Kanehisa M, Furumichi M, Tanabe M, Sato Y, Morishima K (2017) KEGG: New perspectives on genomes, pathways, diseases and drugs. *Nucleic Acids Res* 45:D353–D361.
3. Sorokina M, Stam M, Médigue C, Lespinet O, Vallenet D (2014) Profiling the orphan enzymes. *Biol Direct* 9:10.
4. Shearer AG, Altman T, Rhee CD (2014) Finding sequences for over 270 orphan enzymes. *PLoS One* 9:e97250.
5. Gao J, Ellis LBM, Wackett LP (2010) The University of Minnesota Biocatalysis/Biodegradation Database: Improving public access. *Nucleic Acids Res* 38(Suppl 1): D488–D491.
6. Hatzimanikatis V, et al. (2005) Exploring the diversity of complex metabolic networks. *Bioinformatics* 21:1603–1609.
7. Hatzimanikatis V, Li C, Itonia JA, Broadbelt LJ (2004) Metabolic networks: Enzyme function and metabolite structure. *Curr Opin Struct Biol* 14:300–306.
8. Soh KC, Hatzimanikatis V (2010) DREAMS of metabolism. *Trends Biotechnol* 28: 501–508.
9. Carbonell P, Planson A-G, Fichera D, Faulon J-L (2011) A retrosynthetic biology approach to metabolic pathway design for therapeutic production. *BMC Syst Biol* 5:122.
10. Rodrigo G, Carrera J, Prather KJ, Jaramillo A (2008) DESHARKY: Automatic design of metabolic pathways for optimal cell growth. *Bioinformatics* 24:2554–2556.
11. Cho A, Yun H, Park JH, Lee SY, Park S (2010) Prediction of novel synthetic pathways for the production of desired chemicals. *BMC Syst Biol* 4:35.
12. Yim H, et al. (2011) Metabolic engineering of Escherichia coli for direct production of 1,4-butanediol. *Nat Chem Biol* 7:445–452.
13. Campodonico MA, Andrews BA, Asenjo JA, Palsson BO, Feist AM (2014) Generation of an atlas for commodity chemical production in Escherichia coli and a novel pathway prediction algorithm, GEM-Path. *Metab Eng* 25:140–158.
14. Prather KJ, Martin CH (2008) De novo biosynthetic pathways: Rational design of microbial chemical factories. *Curr Opin Biotechnol* 19:468–474.
15. Delépine B, Duigou T, Carbonell P, Faulon J-L (2017) RetroPath2.0: A retrosynthesis workflow for metabolic engineers. *Metab Eng* 45:158–170.
16. Hadadi N, Hatzimanikatis V (2015) Design of computational retrobiosynthesis tools for the design of de novo synthetic pathways. *Curr Opin Chem Biol* 28:99–104.
17. Carbonell P, Parutto P, Herisson J, Pandit SB, Faulon JL (2014) XTMS: Pathway design in an eXTended metabolic space. *Nucleic Acids Res* 42:W389–W394.
18. Hadadi N, et al. (2014) A computational framework for integration of lipidomics data into metabolic pathways. *Metab Eng* 23:1–8.
19. Rolfsson O, Palsson BØ, Thiele I (2011) The human metabolic reconstruction Recon 1 directs hypotheses of novel human metabolic functions. *BMC Syst Biol* 5:155.
20. Karp PD (2004) Call for an enzyme genomics initiative. *Genome Biol* 5:401.
21. Orth JD, Palsson BO (2010) Systematizing the generation of missing metabolic knowledge. *Biotechnol Bioeng* 107:403–412.
22. Osterman A, Overbeek R (2003) Missing genes in metabolic pathways: A comparative genomics approach. *Curr Opin Chem Biol* 7:238–251.
23. Overbeek R, Fonstein M, D'Souza M, Pusch GD, Maltsev N (1999) The use of gene clusters to infer functional coupling. *Proc Natl Acad Sci USA* 96:2896–2901.
24. Pellegrini M, Marcotte EM, Thompson MJ, Eisenberg D, Yeates TO (1999) Assigning protein functions by comparative genome analysis: Protein phylogenetic profiles. *Proc Natl Acad Sci USA* 96:4285–4288.
25. Chen L, Vitkup D (2006) Predicting genes for orphan metabolic activities using phylogenetic profiles. *Genome Biol* 7:R17.
26. Overbeek R, et al. (2005) The subsystems approach to genome annotation and its use in the project to annotate 1000 genomes. *Nucleic Acids Res* 33:5691–5702.
27. Vallenet D, et al. (2006) MaGe: A microbial genome annotation system supported by synteny results. *Nucleic Acids Res* 34:53–65.
28. Kharchenko P, Chen L, Freund Y, Vitkup D, Church GM (2006) Identifying metabolic enzymes with multiple types of association evidence. *BMC Bioinformatics* 7:177.
29. Yamanishi Y, et al. (2007) Prediction of missing enzyme genes in a bacterial metabolic network. Reconstruction of the lysine-degradation pathway of Pseudomonas aeruginosa. *FEBS J* 274:2262–2273.
30. Chen Y, Mao F, Li G, Xu Y (2011) Genome-wide discovery of missing genes in biological pathways of prokaryotes. *BMC Bioinformatics* 12(Suppl 1):S1.
31. Smith AAT, Belda E, Viari A, Medigue C, Vallenet D (2012) The CanOE strategy: Integrating genomic and metabolic contexts across multiple prokaryote genomes to find candidate genes for orphan enzymes. *PLoS Comput Biol* 8:e1002540.
32. Pearson WR (2013) An introduction to sequence similarity (“homology”) searching. *Curr Protoc Bioinformatics* Chapter 3:Unit3.1.
33. Schnoes AM, Brown SD, Dodevski I, Babbitt PC (2009) Annotation error in public databases: Misannotation of molecular function in enzyme superfamilies. *PLoS Comput Biol* 5:e1000605.
34. Green ML, Karp PD (2007) Using genome-context data to identify specific types of functional associations in pathway/genome databases. *Bioinformatics* 23:i205–i211.
35. Matsuta Y, Ito M, Tohsato Y (2013) ECOH: An enzyme commission number predictor using mutual information and a support vector machine. *Bioinformatics* 29:365–372.
36. Galperin MY, Koonin EV (2012) Divergence and convergence in enzyme evolution. *J Biol Chem* 287:21–28.
37. Ofran Y, Margalit H (2006) Proteins of the same fold and unrelated sequences have similar amino acid composition. *Proteins* 64:275–279.
38. Giri V, Sivakumar TV, Cho KM, Kim TY, Bhaduri A (2015) RxnSim: A tool to compare biochemical reactions. *Bioinformatics* 31:3712–3714.
39. Hu QN, Deng Z, Hu H, Cao DS, Liang YZ (2011) RxnFinder: Biochemical reaction search engines using molecular structures, molecular fragments and reaction similarity. *Bioinformatics* 27:2465–2467.
40. Moriya Y, et al. (2016) Identification of enzyme genes using chemical structure alignments of substrate-product pairs. *J Chem Inf Model* 56:510–516.
41. Hu QN, et al. (2012) Assignment of EC numbers to enzymatic reactions with reaction difference fingerprints. *PLoS One* 7:e52901.
42. Rahman SA, Cuesta SM, Furnham N, Holliday GL, Thornton JM (2014) EC-BLAST: A tool to automatically search and compare enzyme reactions. *Nat Methods* 11: 171–174.
43. Delany J (2011) Daylight Theory Manual, Version 4.9 (DAYLIGHT Chemical Information Systems Inc., Mission Viejo, CA).
44. Rogers DJ, Tanimoto TT (1960) A computer program for classifying plants. *Science* 132:1115–1118.
45. International Union of Biochemistry and Molecular Biology; Webb EC, ed (1992) *Enzyme Nomenclature 1992: Recommendations of the Nomenclature Committee of the International Union of Biochemistry and Molecular Biology on the Nomenclature and Classification of Enzymes* (Academic, San Diego).
46. Hadadi N, Hafner J, Shajkofci A, Zisaki A, Hatzimanikatis V (2016) ATLAS of Biochemistry: A repository of all possible biochemical reactions for synthetic biology and metabolic engineering studies. *ACS Synth Biol* 5:1155–1166.
47. Hadadi N, Hafner J, Soh KC, Hatzimanikatis V (2017) Reconstruction of biological pathways and metabolic networks from in silico labeled metabolites. *Biotechnol J* 12: 1600464.
48. Carbonell P, et al. (2018) Selenzyme: Enzyme selection tool for pathway design. *Bioinformatics* 34:2153–2154.
49. Marmulla R, Safarić B, Markert S, Schweder T, Harder J (2016) Linalool isomerase, a membrane-anchored enzyme in the anaerobic monoterpene degradation in Thauera linaloolentis 47Lol. *BMC Biochem* 17:6.
50. Altschul SF, Gish W, Miller W, Myers EW, Lipman DJ (1990) Basic local alignment search tool. *J Mol Biol* 215:403–410.
51. Briem H, Lessel UF (2002) In vitro and in silico affinity fingerprints: Finding similarities beyond structural classes. *Virtual Screening: An Alternative or Complement to High Throughput Screening?* ed Klebe G (Kluwer Academic Publishers, Dordrecht, The Netherlands), pp 231–244.
52. O'Boyle NM, et al. (2011) Open Babel: An open chemical toolbox. *J Cheminform* 3:33.
53. Weininger D (1988) SMILES, a chemical language and information system. 1. Introduction to methodology and encoding rules. *J Chem Inf Comput Sci* 28:31–36.
54. Leydesdorff L (2008) On the normalization and visualization of author co-citation data: Salton's Cosine versus the Jaccard index. *J Am Soc Inf Sci Technol* 59:77–85.
55. Burgard A, Burk MJ, Osterhout R, Van Dien S, Yim H (2016) Development of a commercial scale process for production of 1,4-butanediol from sugar. *Curr Opin Biotechnol* 42:118–125.
56. Andreozzi S, et al. (2016) Identification of metabolic engineering targets for the enhancement of 1,4-butanediol production in recombinant E. coli using large-scale kinetic models. *Metab Eng* 35:148–159.