

Reconciling Inconsistent Molecular Structures from Biochemical Databases

Casper Asbjørn Eriksen¹, Jakob Lykke Andersen¹, Rolf Fagerberg¹, and Daniel Merkle¹

Department of Mathematics and Computer Science, University of Southern Denmark,
Odense, Denmark

{casbjorn,jlandersen,rolf,daniel}@imada.sdu.dk

Abstract. Information on the structure of molecules, retrieved via biochemical databases, plays a pivotal role in various disciplines, such as metabolomics, systems biology, and drug discovery. However, no such database can be complete, and the chemical structure for a given compound is not necessarily consistent between databases. This paper presents STRUCTRECON, a tool for resolving unique and correct molecular structures from database identifiers. STRUCTRECON traverses the cross-links between database entries in different databases to construct what we call an identifier graph, which offers a more complete view of the total information available on a particular compound across all the databases. In order to reconcile discrepancies between databases, we first present an extensible model for chemical structure which supports multiple independent levels of detail, allowing standardisation of the structure to be applied iteratively. In some cases, our standardisation approach results in multiple structures for a given compound, in which case a random walk-based algorithm is used to select the most likely structure among incompatible alternates. We applied STRUCTRECON to the *EColiCore2* model, resolving a unique chemical structure for 85.11 % of identifiers. STRUCTRECON is open-source and modular, which enables the potential support for more databases in the future.

Keywords: Standardisation · Chemical structure identifiers · Small-molecule databases · Cheminformatics

1 Introduction

As the volume of available biochemical information grows, databases have become indispensable resources for researchers, enabling advances in various fields, including metabolomics, systems biology, and drug discovery. These databases are curated and maintained by different organisations and research groups, each employing their own data collection methods, annotation standards, and quality control procedures. However, as the collective amount of information stored in the databases expands, so does the amount of errors within databases, and in particular, inconsistencies between them [23]. Discrepancies between biochemical

databases pose a significant challenge to researchers performing large-scale analyses, in particular when integrating data from multiple databases [24]. In this work, we focus on incompleteness and inconsistencies in the chemical structures within and between database entries, which can pose a significant problem in applications such as drug discovery, quantitative structure-activity relationship, and atom tracing [25,5].

The entries in each database may contain quantitative information about the compounds, structural information on these compounds, as well as references to related entries in other databases. A starting observation behind our contribution is that the cross-database references can be traversed in order to get a fuller view of the properties of a compound. In the best case, the databases complement each other, making up for the incompleteness of each and allowing the identification a chemical structure of each compound of interest, even if not all of these are contained in any single database.

However, integrating entries from several databases will invariably introduce discrepancies in the chemical structure. In many cases, these discrepancies are simply caused by a difference in the representation of what is intended to be identical chemical structures [14]. Other times, different structural isomers are present under the same name, for example 5-deoxy-D-ribose, which appears in cyclic and linear forms, depending on the database, as depicted in Fig. 1. Over the years, considerable efforts have been directed towards the development of standards and guidelines for chemical structure representation. Structural identifiers such as Standard InChI [8] and Standard SMILES [20,22,21] aim to provide an unambiguous and standardised structural identifier. However, ambiguity is not completely prevented, as sources may wish to denote chemical structures in varying levels of detail, e.g., whether to denote stereochemistry, tautomerism, charge, and more. SMILES makes no distinction whether such features are explicitly represented, while the layered structure of InChI [19] makes it clear which features are to be explicitly represented in some cases.

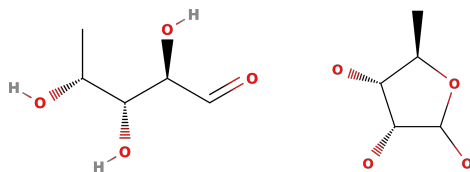


Fig. 1. An example of structural discrepancy between databases: the linear (PubChem, ECMDB, ChEBI) and cyclic (MetaNetX, MetaCyc, KEGG) form of 5-deoxy-D-ribose.

The problem of comparing chemical structures from entries with different notation and level of standardisation is one of the main challenges of this work. It is evident that for biochemical problems, an automatic method for retrieving correct chemical structures, up to some degree of standardisation, is needed. Our goal is to use the combined resources from several databases in order to create

a more complete mapping between database identifiers and chemical structures than any single database provides, while automatically handling discrepancies between these identifiers.

In Section 2, we establish a *model* for representing the chemical structure of compounds, with the goal of being able to describe, and compare across, the various levels of detail to which chemical structures are given by databases. Next, in Section 3, we present STRUCTRECON, a *tool* for programmatically retrieving chemical structures from database identifiers, by traversing database cross-references and using cheminformatics methods for analysing, comparing, and standardising structural representations based on the model developed in Section 2. Finally, in section 4, we apply the tool to a set of compounds established by genome sequencing of *E. coli* and analyse the resulting network of identifiers and structural representations.

2 Multi-Level Modelling of Chemical Structures

In this section, we introduce a model for chemical structures which allows representation at multiple levels of detail and formalises the standardisation functions which transform structures between these levels. Compared to established models, such as SMILES and InChI, this model places a particular focus on extensibility, formal specification, and standardisation of structures.

We call the levels of detail *features*. The seven features used throughout this work will be introduced one-by one as the necessary theory is established. We categorise identifiers into two classes: *structural identifiers*, which directly encode a chemical structure, and from which the structure can be recovered algorithmically (e.g. InChI, SMILES), and *symbolic identifiers*, which are generated more or less arbitrarily, and do not carry direct meaning, but reference an entry in the corresponding database (e.g. PubChem CID, BiGG ID).

Depending on the application, the exact definition of a chemical structure may vary. We wish to model the connectivity of atoms in a molecule, and optionally stereo-chemical information, while the full spatial information is not taken into account. For this application, the classical method in cheminformatics is the graph-based approach: a structure is, in its most basic form, an unlabelled graph $G = (V, E)$, in which atoms are represented by vertices, and (covalent) bonds are represented by edges. A graph-based structure is linearly encoded by established structural identifiers, such as SMILES and InChI; even the systematic IUPAC name may be described as a graph-based identifier [10]. We may wish to describe additional features, e.g., charge, isotopic labelling, stereochemistry, or tautomerism. In broad terms, a chemical feature is a class of information about a structure which we may or may not wish to account for when modelling, depending on the input and application.

We will establish a unified model for graph-based chemical structures. This provides a consistent view of chemical structure, regardless of the features represented, to which structural identifiers can be mapped. The underlying simple graph structure is, for our purposes, always assumed to be present, but even

basic information such as the chemical element of each atom and the order of covalent bonds, are considered optional features. First, we formally define the notion of a feature.

Definition 1 (Feature). *A feature Φ is a pair $\Phi = (\Phi_V, \Phi_E)$, where Φ_V and Φ_E are sets of possible values for the attribute on atoms and bonds respectively, each of which must contain a special ‘nil’-element, ϵ , indicating that the value is unspecified or not applicable.*

We will define and apply seven such features in this work. Starting with the most essential, the element of each atom can be expressed as a feature \mathbf{E} , with

$$\mathbf{E} := (\{\epsilon, \text{H}, \text{He}, \text{Li}, \dots\}, \{\epsilon\}).$$

For a given structure, each vertex will either be assigned ϵ , indicating that no indication is given as to the element of that atom, or it will be assigned a specific element. Edges can only be assigned ϵ , indicating that this feature assigns no attribute to edges. Similarly, bond types are expressed as the feature

$$\mathbf{B} := (\{\epsilon\}, \{\epsilon, -, =, \equiv\}).$$

This feature can be expanded to also indicate other bond types, such as aromatic or ionic, if needed. The isotope of each atom can be stated as

$$\mathbf{I} := (\{\epsilon\} \cup \mathbb{N}, \{\epsilon\}),$$

where the vertex attribute indicates the atomic weight of each atom. Before describing the remaining features, we need a precise definition of chemical structure. Combining the sets of values for all features, we can define the overall feature space, which is needed for a formal definition of a chemical structure.

Definition 2 (Feature space). *Given a set of features $\{\Phi_1, \dots, \Phi_n\}$, the feature space for vertices and edges, \mathcal{F}_V and \mathcal{F}_E , respectively, is the combined attribute space of the features, where $\mathcal{F}_V = \Phi_{1V} \times \dots \times \Phi_{nV}$ and $\mathcal{F}_E = \Phi_{1E} \times \dots \times \Phi_{nE}$.*

Definition 3 (Chemical structure). *A chemical structure is an undirected graph $G = (V, E, A_V, A_E)$, where V is the set of atoms, E is the set of covalent bonds, and $A_V: V \rightarrow \mathcal{F}_V$ (resp. $A_E: E \rightarrow \mathcal{F}_E$) is an attribute function, assigning to each vertex (resp. edge) a value for each feature in the feature space.*

Let \mathcal{G} be the set of all such chemical structures. Simultaneously working in several levels of detail, i.e., features, naturally raises the problem of how to compare equivalence of structures between different sets of features. For this, we introduce, for each feature, a *standardisation function*.

Definition 4 (Standardisation function). *Given a feature Φ , the standardisation function w.r.t. Φ is $S_\Phi: \mathcal{G} \rightarrow \mathcal{G}$. The function is required to be idempotent, and all vertices and edges of the resulting structure should have ϵ as value for feature Φ in the attribute function.*

We extend the definition to sequences of features, resulting in a composition of standardisations, i.e., $S_{\Phi_1 \dots \Phi_n} = S_{\Phi_n} \circ \dots \circ S_{\Phi_1}$.

For a structure G , the image $S_{\Phi}(G)$ is the corresponding standardised structure, which does not contain any information about feature Φ . For a sequence of features Φ_1, \dots, Φ_n , we define $\mathcal{G}_{\Phi_1 \dots \Phi_n} \subseteq \mathcal{G}$ as the set of structures which are standardised according to the of features. That is, all applicable atom and bond attributes are ϵ these features, and they are their own image in the standardisation function:

$$G \in \mathcal{G}_{\Phi_1 \dots \Phi_n} \iff S_{\Phi_1 \dots \Phi_n}(G) = G.$$

For the features described so far, **E**, **B**, and **I**, the *trivial standardisation function* is sufficient. This function simply erases the attributes by setting them to ϵ . In some cases, the trivial standardisation function is not sufficient. For example, we define the charge feature as

$$\mathbf{C} := (\{\epsilon\} \cup \mathbb{Z}, \{\epsilon\}),$$

where the vertex attribute indicates the charge of the atom. The process of standardising the charge of a given molecule may be limited by chemical constraints, such as is the case in RDKit [16] and InChI [19]. The function may, among other modifications, add and/or remove hydrogen atoms, trying to remove charges from each atom in a chemically valid way. We will not further discuss the intricacies of this operation, as it is implementation-dependant.

For some features, the definition even depends entirely on the standardisation function. E.g., the fragment feature [14]:

$$\mathbf{F} := (\{\epsilon\}, \{\epsilon\}).$$

In this case, no information is explicitly encoded in the graph, but we still want a way to standardise w.r.t. fragments. We define $S_{\mathbf{F}}(G)$ to be the largest connected component of G (measured by the number of non-hydrogen atoms), with an implementation-specific method for breaking ties based on other attributes of each connected component.

Tautomerism is similarly difficult to define due to the complexities involved in determining the tautomers of any given structure. Let

$$\mathbf{T} := (\{\epsilon\}, \{\epsilon\}).$$

Again, we have no features, instead relying on the standardisation function: assume a chemical oracle, which given a structure G , returns the set of all tautomeric structures, according to some definition. Then, let a deterministic method choose a canonical representative from among these structures. The standardised structure $S_{\mathbf{T}}(G)$ is this canonical tautomer.

For stereochemistry, there are multiple methods for encoding information about local geometry at the vertex and edge level [2,1,15] Any method for encoding such information can be used to generate the feature **S**. The trivial standardisation function will in most cases be sufficient for standardising structures with respect to stereochemistry.

Finally, we define equivalence of chemical structures: With the notation $G_1 =_{\Phi_1 \dots \Phi_n} G_2$, we denote that the structures G_1 and G_2 are equivalent up to standardisation of Φ_1, \dots, Φ_n , i.e., they are equal when the features Φ_1, \dots, Φ_n are not considered. Formally

$$G_1 =_{\Phi_1 \dots \Phi_n} G_2 \iff S_{\Phi_1, \dots, \Phi_n}(G_1) = S_{\Phi_1, \dots, \Phi_n}(G_2).$$

As an example, consider the structures for methanol and methoxide, where $\text{CH}_3\text{OH} \neq_{\text{FICTS}} \text{CH}_3\text{O}^-$, but $\text{CH}_3\text{OH} =_{\text{FITS}} \text{CH}_3\text{O}^-$.

We have now described the features which are considered in this contribution, based on the FICTS features [18]: (**E**) elements, (**B**) bonds, (**F**) fragments, (**I**) isotopes, (**C**) charge, (**T**) tautomerism, and (**S**) stereoisomerism. It should be noted that some of these features depend upon each other, e.g., it would not make sense to specify the isotope of an atom without also specifying the element. The standardisation functions can not be expected commute with each other in general. For this reason, STRUCTRECON needs a defined order in which the standardisation functions will be applied.

3 Algorithms and Implementation

In this section, we describe the ideas and algorithms of STRUCTRECON, based on the model developed in Section 2.

Data sources We used six sources of data: BiGG [13], ChEBI [4], the E. Coli Metabolome Database (ECMDB) [7,17], KEGG [11], MetaNetX [6], and PubChem [12]. These were selected based on their programmatic accessibility and relevance to metabolic modelling. STRUCTRECON is modular, making it easy to add more data sources in the future.

MetaNetX uses a versioning system in which entries present in one version are not necessarily present in another. Inter-database references do not typically specify which version of MetaNetX is referenced, but MetaNetX keeps a record of deprecated entries. Following the chain of deprecations, it is possible to obtain the newest entries corresponding to any given entry. The deprecation relationship can both split and merge entries in between versions.

Construction of the Identifier Graph The interconnected nature of identifiers within and between databases is represented as a directed graph, called the *identifier graph*. In the identifier graph, vertices correspond to identifiers, while arcs represent relationships between these. Each vertex contains as attributes, the type of ID (e.g. PubChem CID, BiGG ID, SMILES) called the *identifier class*, as well as the actual ID. Each edge is annotated with the source database.

A number of *procedures* are specified, each being a subroutine which takes identifiers as input, and finds associated identifiers in chemical databases. By executing the procedures in an order which seeks to minimise overhead, the identifier graph is built iteratively, starting with the *input vertices* which are directly obtained as input. The resulting graph will contain symbolic identifiers as well as the structural identifiers as they appear in the respective databases. For an example of a complete identifier graph, refer to Fig. 2.

Structure Standardisation When a structural identifier (SMILES, InChI) is added to the identifier graph, it is first converted to an internal graph-based representation, in accordance with Def. 3. From this point, we assume that the atom (**A**) and bond (**B**) features are always implicitly represented, and will therefore refer to the remaining features as **FICTS**. As the standardisation functions are not expected to commute, in general, we enforce a particular order on the features defined by the user, by default **F, I, C, T, S**. This was chosen as the default ordering, as it produces the greatest number of uniquely resolved structures which a lesser degree of standardisation.

We assign to each structural identifier G , an attribute specifying for which features a structure is standardised. For features **F**, **I**, and **S**, it is simple to guarantee that it is standardised by inspecting the structure. Checking **F** is simply examining the connectivity of the graph, and **I**, **S**, inspecting for all vertices and edges whether they have the equivalent of ϵ as attribute. For **C** and **T**, we check whether a structure is standardised by applying the respective standardisation functions.

When the links to new databases is exhausted, in many cases, there will be multiple different structural identifiers associated with each compound. We aim to achieve a unified representation of the compounds by iteratively applying the standardisation functions: For any structure G , which is not fully standardised, that is, $G \notin \mathcal{G}_{\text{FICTS}}$, let Φ_k be the first feature in which G is not standardised according to the feature ordering. That is, $S_{\Phi_1 \dots \Phi_{k-1}}(G) = G$, but $S_{\Phi_1 \dots \Phi_k}(G) \neq G$. Then produce $G' = S_{\Phi_k}(G)$, adding this new structure to the identifier graph, with an arc (G, G') . If G' is equivalent to an existing structure H , then no new vertices are created, but the arc (G, H) is added instead. The standardisation process is visualised in the output of STRUCTRECON (Fig. 2) in which the blue nodes represent structures, and purple nodes represent maximally standardised structures.

Structure Selection While the identifier graph is a general digraph, considering only the structural identifiers yields a forest of in-trees, as standardisation functions are many-to-one, and may therefore merge structures, but never split. For each input, the transitive and reflexive closure of the forest of in-trees imposes a partial order on the reachable structures. In this partial order, structure G_1 precedes G_2 if G_2 is a more standardised identifier, reachable from G_1 by applying standardisation functions. Maximal elements are fully standardised structures, $G \in \mathcal{G}_{\text{FICTS}}$.

For each input identifier i , STRUCTRECON should resolve the input to a single structure. The vertices reachable from i represent symbolic and structural identifiers which are related to i through database links, as well as structural identifiers which can be derived from these by standardisation. Denote by $\mathcal{S}(i)$ the partial order of structures reachable from i .

If $\mathcal{S}(i)$ has one maximal (greatest) element, then we say that $\mathcal{S}(i)$ is resolved and *consistent* — all sources for the compound can agree on a structure, at least up to the highest degree of standardisation. That is, $G =_{\text{FICTS}} G'$ for any pair of structures G, G' in the reachable database entries. If there are multiple

maximal elements, then the sources cannot agree, and we call $\mathcal{S}(i)$ *inconsistent*. If $\mathcal{S}(i)$ contains no structures at all, then the compound is *unresolved*. In the case where $\mathcal{S}(i)$ is consistent, we want to select the most specific element on which all sources agree. In the partial order, this is the supremum of the set of structures which were found directly in the databases.

If $\mathcal{S}(i)$ is inconsistent, resolving i requires a choice between the maximal structures. This choice can be made automatically by a scoring algorithm which, for each input, assigns a confidence score to all vertices reachable from the input vertex. The scoring algorithm essentially computes the probability that a random walk in the identifier graph, starting at the input vertex, will arrive at each vertex. The algorithm is based on PageRank [3], with the key differences that the initial probability distribution is 1 for the input vertex, and 0 for all other vertices, and that sink vertices only loop back to the input vertex, rather than all vertices. After assigning a confidence score to each vertex, the confidence scores of each maximally standardised structure is evaluated. The *confidence ratio* is computed, the confidence of the second-most likely structure over the confidence of the most likely structure. If this value is below a given threshold (0.5 by default), then the vertex with the highest confidence score is automatically selected. Otherwise, no structure can be chosen with high enough confidence, and the input is marked for manual disambiguation.

Implementation STRUCTRECON was implemented in Python, and is available at <https://github.com/casbjorn/structrecon>. The accompanying web interface can be accessed at <https://cheminf.imada.sdu.dk/structrecon>. Our model for chemical structure is implemented using RDKit [16], which furthermore provides functions for uncharging structures and computing the canonical tautomer.

4 Results

The tool was tested on the metabolic network model *EColiCore2* [9]. The model contains 2138 identified compounds, with associated BiGG IDs. We chose this dataset for evaluation, as the selected databases have a particular focus on biochemistry, and because of the well-established nature of the *E. coli* genome.

Of the 2138 inputs, 136 (6.36 %) were identified as macromolecules, based on string-matching BiGG IDs and names found in databases to an incomplete list of substrings associated with macromolecules, such as "tRNA" and the names of various proteins and enzymes. We consider the handling of macromolecules to be out of the scope of this work. In the identifier graph, an average of 31.70 vertices are reachable from each input vertex. Of the non-macromolecule inputs, 1459 (72.88 %) resolved to exactly one structure up to maximal **FICTS** standardisation, while only 492 (24.58 %) had only one structure up to **FICT** standardisation. Of the non-macromolecule inputs, 57 (2.85 %) yielded no structure at all. Examples of this category includes bis-molybdopteringuaninedinucleotide (BiGG: M_bmocogdp), Hexadecanoyl-phosphate(n-C16:1) (BiGG: M_hdceap), and 2-tetradec-7-enoyl-sn-glycerol-3-phosphate (BiGG: M_2tdec7eg3p).

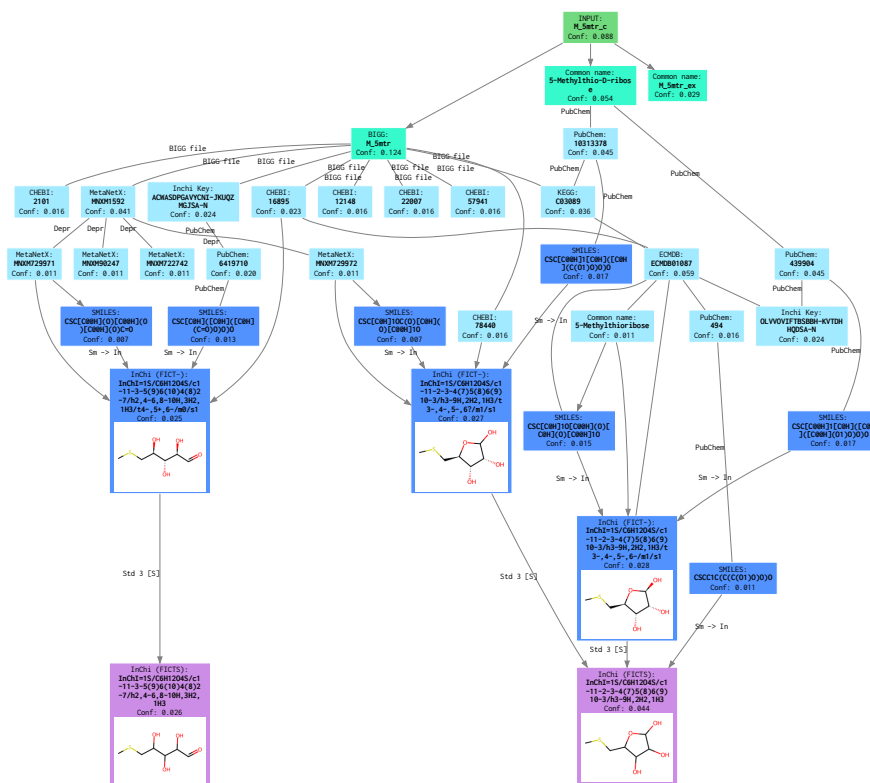


Fig. 2. The identifier graph generated by the BiGG ID *M_5mtr_c*. Each vertex displays the type of identifier, the identifier itself, and the confidence assigned by the scoring algorithm. For structures, the set of features in which the structure is standardised is also displayed, along with a graphical representation. The green vertex is the input vertex. The turquoise vertices are symbolic identifiers found directly within the input file, in this case the *EColiCore2* model. The light blue vertices are other symbolic identifiers. The dark blue vertices represent structural identifiers, either found in databases, or obtained by standardisation. The violet vertices represent maximally standardised structures. Arcs with no direction are shorthand for one arc in each direction.

A total of 486 inputs (24.28%) yielded multiple maximally standardised structures, and needed to be disambiguated based on the *confidence ratio*. In our experimentation, we found 0.5 to be a reasonable threshold, meaning that we select the structure with the highest confidence if it has at least twice the confidence of any other structure. With a threshold of 0.5, an additional 245 inputs were uniquely resolved, for a total of 1704 consistent inputs (85.11%), leaving 241 compounds for manual disambiguation. The effect of different choices of confidence ratio threshold is displayed in Fig. 3.

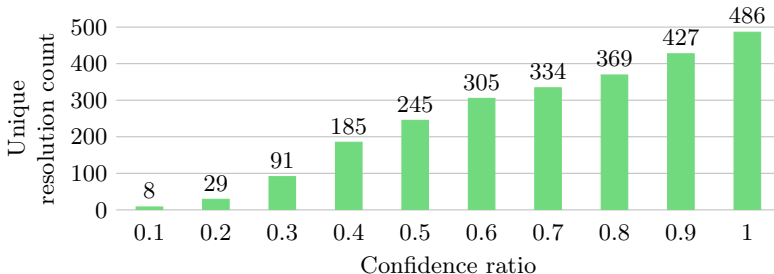


Fig. 3. For several choices of confidence ratio threshold (the confidence of the second-most likely structure over the confidence of the most likely structure), shows the number of inputs, out of 486, which resolve to a unique structure in the ECC2 model. STRUCTRECON uses a default threshold of 0.5. Setting the threshold to 0.0 would mean only choosing a structure if no alternatives exist, while a threshold of 1.0 results in picking one of the structures with the highest confidence arbitrarily.

We will proceed to describe some concrete examples of identifier graphs which serve to demonstrate both the problem of database inconsistency and the solution provided by STRUCTRECON. One example is 5-methylthio-D-ribose. The associated identifier graph is displayed in Fig. 2. Database interconnections do not necessarily make distinctions between this compound and S-methyl-5-thio-D-ribose. The confidence ratio between these two maximised structures is 0.59, indicating a relatively high degree of interconnection between the associated database entries. While the correct structure has the highest confidence score, the default threshold of 0.5 would mark this discrepancy for manual disambiguation.

Unexpectedly, the simplest and most prevalent molecules turns out to be inconsistent, but easy to reconcile based on the confidence ratio. A good example is water, as displayed in Appendix A. The conventional structure H₂O is found in a multitude of databases, however, the ChEBI identifier 29356 (oxide(2-)) is associated with the generic BiGG identifier for water through the BiGG database. However, as this is the only connection, that structure is assigned a smaller confidence score than the conventional structure by the scoring algorithm. This graph therefore has a low confidence ratio of 0.07, representing a high degree of support of the conventional structure, which is chosen by STRUCTRECON.

5 Conclusion

In this work, we propose a model for chemical structure, which supports multiple levels of standardisation. Based on this model, we present STRUCTRECON, a tool which identifies and reconciles the chemical structure of compounds based on the traversal of interconnections between biochemical databases. We applied the tool to *EColiCore2*, a metabolic model of *E. coli*. In 85.11 % of cases, a chemical structure could be uniquely identified with reasonable confidence, demonstrat-

ing that STRUCTRECON can be a valuable tool for structure-based approaches in bioinformatics and related fields. STRUCTRECON is open-source and developed with modularity in mind, making integration of additional databases and procedures possible.

Acknowledgements This work was supported by the Novo Nordisk Foundation (grant NNF21OC0066551, ‘MATOMIC’).

References

1. Akutsu, T.: A new method of computer representation of stereochemistry. transforming a stereochemical structure into a graph. *Journal of Chemical Information and Computer Sciences* **31**(3) (1991). <https://doi.org/10.1021/ci00003a008>
2. Andersen, J.L., Flamm, C., Merkle, D., Stadler, P.F.: Chemical Graph Transformation with Stereo-Information. In: de Lara, J., Plump, D. (eds.) *Graph Transformation*. pp. 54–69. *Lecture Notes in Computer Science*, Cham (2017). https://doi.org/10.1007/978-3-319-61470-0_4
3. Brin, S., Page, L.: The anatomy of a large-scale hypertextual web search engine. *Computer Networks and ISDN Systems* **30**(1) (1998). [https://doi.org/10.1016/S0169-7552\(98\)00110-X](https://doi.org/10.1016/S0169-7552(98)00110-X)
4. Degtyarenko, K., de Matos, P., Ennis, M., Hastings, J., Zbinden, M., McNaught, A., Alcántara, R., Darsow, M., Guedj, M., Ashburner, M.: ChEBI: a database and ontology for chemical entities of biological interest. *Nucleic Acids Research* **36**(Database issue), D344–D350 (Jan 2008). <https://doi.org/10.1093/nar/gkm791>
5. Fourches, D., Muratov, E., Tropsha, A.: Trust, But Verify: On the Importance of Chemical Structure Curation in Cheminformatics and QSAR Modeling Research. *Journal of Chemical Information and Modeling* **50**(7), 1189–1204 (Jul 2010). <https://doi.org/10.1021/ci100176x>
6. Ganter, M., Bernard, T., Moretti, S., Stelling, J., Pagni, M.: MetaNetX.org: a website and repository for accessing, analysing and manipulating metabolic networks. *Bioinformatics* **29**(6), 815–816 (Mar 2013). <https://doi.org/10.1093/bioinformatics/btt036>
7. Guo, A.C., Jewison, T., Wilson, M., Liu, Y., Knox, C., Djoumbou, Y., Lo, P., Mandal, R., Krishnamurthy, R., Wishart, D.S.: Ecmdb: the e. coli metabolome database. *Nucleic Acids Research* **41**(Database issue), D625–630 (Jan 2013). <https://doi.org/10.1093/nar/gks992>
8. Heller, S.R., McNaught, A., Pletnev, I., Stein, S., Tchekhovskoi, D.: InChI, the IUPAC International Chemical Identifier. *Journal of Cheminformatics* **7**(1), 1–34 (Dec 2015). <https://doi.org/10.1186/s13321-015-0068-4>
9. Hädicke, O., Klamt, S.: Ecolicore2: a reference network model of the central metabolism of escherichia coli and relationships to its genome-scale parent model. *Scientific Reports* **7**(11) (2017). <https://doi.org/10.1038/srep39647>
10. International Union of Pure and Applied Chemistry Commission on the Nomenclature of Organic Chemistry, Klesney, S.P.: *Nomenclature of Organic Chemistry* (1979)
11. Kanehisa, M., Furumichi, M., Tanabe, M., Sato, Y., Morishima, K.: KEGG: new perspectives on genomes, pathways, diseases and drugs. *Nucleic Acids Research* **45**(D1), D353–D361 (Jan 2017). <https://doi.org/10.1093/nar/gkw1092>

12. Kim, S., Chen, J., Cheng, T., Gindulyte, A., He, J., He, S., Li, Q., Shoemaker, B.A., Thiessen, P.A., Yu, B., Zaslavsky, L., Zhang, J., Bolton, E.E.: PubChem 2023 update. *Nucleic Acids Research* **51**(D1), D1373–D1380 (Jan 2023). <https://doi.org/10.1093/nar/gkac956>
13. King, Z.A., Lu, J., Dräger, A., Miller, P., Federowicz, S., Lerman, J.A., Ebrahim, A., Palsson, B.O., Lewis, N.E.: BiGG Models: A platform for integrating, standardizing and sharing genome-scale models. *Nucleic Acids Research* **44**(D1), D515–D522 (Jan 2016). <https://doi.org/10.1093/nar/gkv1049>
14. Muresan, S., Sitzmann, M., Southan, C.: Mapping Between Databases of Compounds and Protein Targets. *Methods in molecular biology* (Clifton, N.J.) **910**, 145–164 (2012). https://doi.org/10.1007/978-1-61779-965-5_8
15. Petrarca, A.E., Lynch, M.F., Rush, J.E.: A method for generating unique computer structural representations of stereoisomers. *Journal of Chemical Documentation* **7**(3) (1967). <https://doi.org/10.1021/c160026a008>
16. RDKit: Open-source cheminformatics software, <https://www.rdkit.org/>
17. Sajed, T., Marcu, A., Ramirez, M., Pon, A., Guo, A.C., Knox, C., Wilson, M., Grant, J.R., Djoumbou, Y., Wishart, D.S.: Ecmdb 2.0: A richer resource for understanding the biochemistry of e. coli. *Nucleic Acids Research* **44**(D1), D495–501 (Jan 2016). <https://doi.org/10.1093/nar/gkv1060>
18. Sitzmann, M., Filippov, I., Nicklaus, M.: Internet resources integrating many small-molecule databases1. *SAR and QSAR in Environmental Research* **19**(1-2), 1–9 (Jan 2008). <https://doi.org/10.1080/10629360701843540>
19. Stein, S.E., Heller, S.R., Tchekhovskoi, D.V.: The IUPAC Chemical Identifier – Technical Manual
20. Weininger, D.: SMILES, a chemical language and information system. 1. Introduction to methodology and encoding rules. *Journal of Chemical Information and Computer Sciences* **28**(1), 31–36 (Feb 1988). <https://doi.org/10.1021/ci00057a005>
21. Weininger, D.: SMILES. 3. DEPICT. Graphical depiction of chemical structures. *Journal of Chemical Information and Computer Sciences* **30**(3), 237–243 (Aug 1990). <https://doi.org/10.1021/ci00067a005>
22. Weininger, D., Weininger, A., Weininger, J.L.: SMILES. 2. Algorithm for generation of unique SMILES notation. *Journal of Chemical Information and Computer Sciences* **29**(2), 97–101 (May 1989). <https://doi.org/10.1021/ci00062a008>
23. Williams, A.J., Ekins, S.: A quality alert and call for improved curation of public chemistry databases. *Drug Discovery Today* **16**(17), 747–750 (Sep 2011). <https://doi.org/10.1016/j.drudis.2011.07.007>
24. Williams, A.J., Ekins, S., Tkachenko, V.: Towards a gold standard: regarding quality in public domain chemistry databases and approaches to improving the situation. *Drug Discovery Today* **17**(13), 685–701 (Jul 2012). <https://doi.org/10.1016/j.drudis.2012.02.013>
25. Young, D., Martin, T., Venkatapathy, R., Harten, P.: Are the Chemical Structures in Your QSAR Correct? *QSAR & Combinatorial Science* **27**(11-12), 1337–1345 (2008). <https://doi.org/10.1002/qsar.200810084>

Appendix A

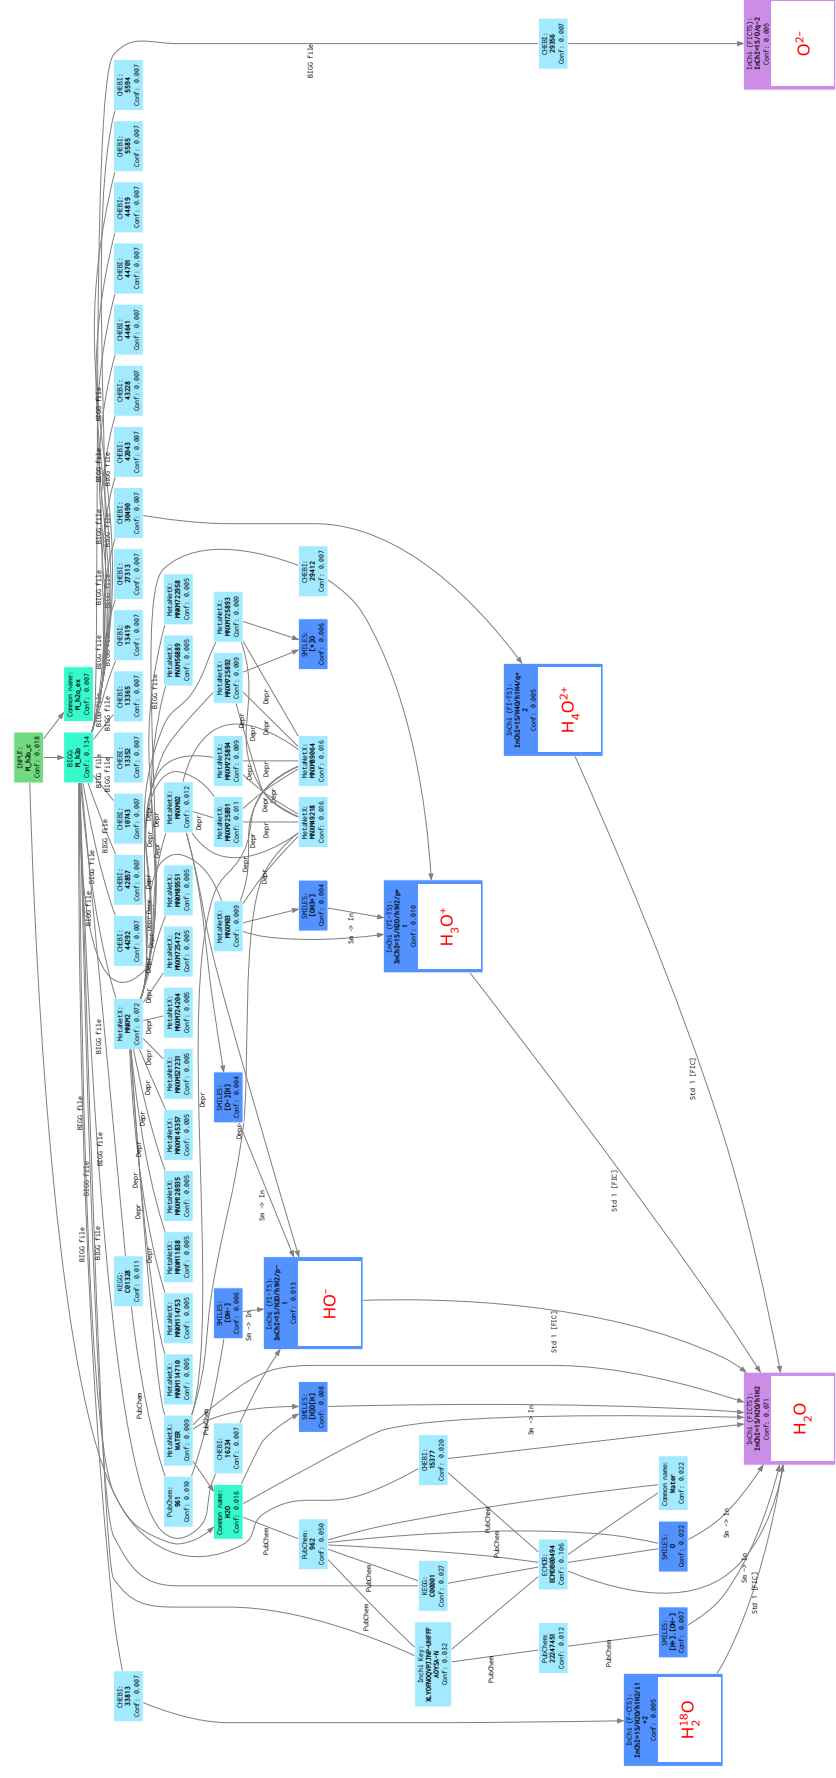


Fig. 4. The identifier graph generated by the BiGG ID M_h20_c. For a description of how to read this graph, see the caption of Fig. 2. This figure clearly demonstrates the impact of the scoring algorithm, as it chooses the conventional structure, H₂O, at a confidence ratio of 0.07.