

Representing Catalytic Mechanisms with Rule Composition

Jakob L. Andersen, Rolf Fagerberg, Christoph Flamm, Walter Fontana, Juri Kolčák,*
Christophe V.F.P. Laurent, Daniel Merkle,* and Nikolai Nøjgaard



Cite This: *J. Chem. Inf. Model.* 2022, 62, 5513–5524



Read Online

ACCESS |



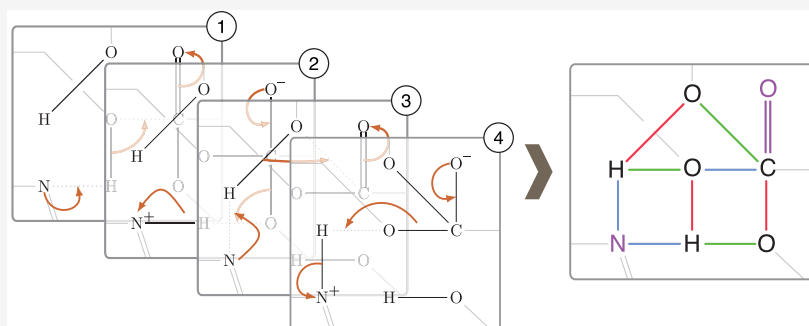
Metrics & More



Article Recommendations



Supporting Information



ABSTRACT: An “imaginary transition structure” overlays the molecular graphs of the educt and product sides of an elementary chemical reaction in a single graph to highlight the changes in bond structure. We generalize this idea to reactions with complex mechanisms in a formally rigorous approach based on composing arrow-pushing steps represented as graph-transformation rules to construct an overall composite rule and a derived transition structure. This transition structure retains information about transient bond changes that are invisible at the overall level and can be constructed automatically from an existing database of detailed enzymatic mechanisms. We use the construction to (i) illuminate the distribution of catalytic action across enzymes and substrates and (ii) to search in a large database for reactions of known or unknown mechanisms that are compatible with the mechanism captured by the constructed composite rule.

INTRODUCTION

A chemical reaction is often viewed as occurring in a sequence of steps considered to be elementary at some level of abstraction. One such level is the concerted net movement of one or two electrons from a center high in electronic density to a center low in electronic density, as codified by “arrow pushing”.¹ A mechanism is then a sequence of elementary steps—or steps, for short—that reach a stable configuration. At a higher level of abstraction, one might view a complete reaction as an “elementary” step and a synthesis path or a metabolic pathway as a mechanism.

A mechanism explains the chemistry of a reaction by providing a causal account of how (and why) a particular outcome is attained. Cyclic mechanisms, such as catalysis, are of particular interest in biology and industry. Since the catalyst is recycled, any molecular change it undergoes in one step of the mechanism must be undone in subsequent steps. Much of the chemistry may then consist of transient changes that remain invisible at the granularity of the overall reaction, which adds up all the steps. This is the granularity at which many enzymatic reactions are expressed. Our aim in this paper is to develop a representation of a catalyzed multi-step reaction that is compact like an overall reaction yet richer in that it captures transient modifications.

Although fairly intuitive to grasp, this representation must be made formal. Formality provides it with clear computational semantics for automatic construction and deployment. We approach this task within the framework of graph transformation.² It is standard practice to draw a molecular structure formula as an undirected graph, with vertices labeled by atom types and edges by bond types. Since a reaction is the modification of educt graphs into product graphs, methods for expressing graph transformations³ provide a natural representation of chemistry. A graph transformation is, roughly, a rule $L \rightarrow R$, with L and R denoting structural patterns. Any molecular structure graph that matches L qualifies for the transformation, which consists in replacing the matching part with R . This is akin to what chemists call a reaction template. The difference is that a graph transformation must specify which atoms in L correspond to which atoms in R , that is, the atom map. It is

Received: April 13, 2022

Published: November 3, 2022



because of the atom-map requirement that a rule captures an aspect of the mechanism not provided by a reaction template. This requirement is also necessary for the notion of graph transformation to be expressed in a mathematically rigorous manner.

To the best of our knowledge, currently, the only open-source graph-transformation platform for chemistry based on an atom-level graph model of molecules is MØD.⁴ Given a collection of reactants, MØD identifies the graphs that match a pattern specified by a rule and returns the products that result from applying the rule. In this way, MØD can apply a chemistry codified as a set of rules to an initial set of reactants and iteratively determine the implied reaction network to any depth as computational resources permit.

Unfortunately, atom maps of chemical reactions are not provided as often as one would hope. However, elementary steps or reactions are typically expressed in terms of “arrow-pushing” diagrams, which clearly imply the atom map. Such steps can be translated into graph-transformation rules. A multi-step reaction, such as an enzyme-catalyzed reaction, can then be represented as the sequential application of step-specific rules. The main contribution of the work presented here consists in expressing this sequence succinctly as a single rule obtained by the formal composition of all step-specific rules. It turns out that an equivalent representation of this composite rule is an overlay of so-called imaginary transition structures (ITSs) for each elementary step. The notion of an ITS was introduced by Fujita.⁵ An ITS is a structural formula in which the right-hand side of a reaction is superposed on the left, yielding a classification of each bond according to whether it is present only in the educts, only in the products, or in both. The overlay graph (OG) introduced in the present paper naturally extends the concept of ITS with additional classes of transient changes. Tracking transient changes is particularly important when dealing with catalytic mechanisms since the catalyst may, and indeed often does, swap atoms with the substrate(s).

We believe the OG to be a useful conceptual and computational tool for describing, analyzing, searching, and designing catalyzed reactions. Some use cases in support of this belief are provided. In the first application, we exploit arrow-pushing information in the Mechanism and Catalytic Site Atlas⁶ (M-CSA) to automatically construct OGs of hydrolase reactions listed in the database. These OGs point at a spectrum of catalytic entanglements of enzymes and substrates, de-emphasizing the notion of a singular catalyst in favor of a collection of catalytic sites that can be distributed across the enzymes and substrates. In the second application, we deploy composite rules to automatically search the Rhea database⁷ for reactions of known or unknown mechanisms (i.e., no atom map) that are compatible with the mechanism captured by a composite rule.

Since the graph-transformation approach can be heavy for those unfamiliar with it, we structured the main text of the paper to be largely devoid of formalism.

2. METHODS

In this section, we convey the idea of an OG by sketching its construction in an informal manner. A rigorous specification, capable of guiding software implementation, is provided in the Supporting Information at two levels of sophistication—one that is mathematically precise yet still close to chemistry, and another that is very formal by deploying the language of

category theory, which is the lingua franca for those in the mathematical branch of the graph-transformation business.

To be meaningfully informal, however, we still need to clarify our vocabulary. The term “rule” is typically used by chemists interchangeably with the notion of a reaction template. The idea of a reaction template, $L \rightarrow R$, is to only write those parts, L , of molecular structures that are required for a reaction and to specify that L is to be rearranged into R . Some parts of L might be atom groups that enable chemical transformation (by shaping electronic density or playing a steric role) but are not themselves modified. Chemists represent molecular structures as graphs, and a pattern is a molecular graph in which some parts are omitted. In the literature, this omission is typically flagged with some kind of a remainder symbol. While a remainder symbol like $-R$ is a useful visual cue, it is a distraction here; we simply leave off the parts that we do not care about. To determine whether a template can be applied, the reactants must match the pattern L . By rearranging the matching parts in the reactants according to the right side R of the template, we obtain a reaction between fully specified molecular structures. Templates are often applied to a collection of molecules by selecting a combination of reactants that match. We use the term “mixture” as a shorthand for “a collection of molecules”. A mixture can be very large, like the contents of a test tube, but it can equally refer to a more abstract situation comprising a few reactants that are all used up in a single template application.

The key point to note is that this notion of the template does not necessarily capture the mechanism of a reaction because it does not require knowing which atoms on the left correspond to which on the right. An atom map is necessary for specifying a physical process of how parts are rearranged beyond solely stating their rearrangement. (The atom map is insufficient for specifying the full mechanism since it does not fix the temporal order of physical changes.) While the atom map is not needed for many chemical applications, it plays a crucial role in this paper. Throughout this contribution, we therefore reserve the term “rule” for a reaction template augmented with an atom map. A rule is formally defined in the Supporting Information (Section A). On a less formal level, drawings suffice. The top row of Figure 1A illustrates a rule, and the bottom row represents its application to a mixture, consisting in this case of a single molecule. The arrow labeled m indicates a matching.

Figure 1A points to a subtlety regarding patterns that we need to clarify. The omission of structural parts in forming a pattern means that we do not care which other parts take their place when deciding whether a molecule matches the pattern. However, the omission of a bond between two particular atoms is ambiguous as it could mean that the bond is optional or that its absence is required. For example, the rule in Figure 1A creates a bond between an H and a C atom and thus requires them not to be bound at the outset. We use a dotted line between two specific atoms to assert the required absence of a bond between them (similarly, as also illustrated in Figure 1A, breaking a bond between two specific atoms results in a product pattern R , in which these atoms are doubtlessly not bound to each other, a fact flagged by a dotted line). Dotted lines never appear in a mixture where molecules are fully specified.

Fujita’s ITS, shown in Figure 1B, is but an alternative notation for a rule. Instead of using an arrow separating the




Figure 1. (A): Illustrative example of a rule describing a keto–enol tautomerization, applied to 2-hydroxyacrylaldehyde. The dotted lines represent the non-bond constraint (see text) that the corresponding atoms not be bound to each other. For example, on the educt side, the C and H atoms must not be bound to each other since the rule creates that bond. Given the informal presentation in this section, we only insinuate the existence of an atom map between the educt and product molecules by using a spatial layout, suggesting which atoms on the left of a rule correspond to which atoms on the right. (B): ITS⁵ of the reaction in panel A. The ITS is an overlay of the R and L patterns of a rule. In a rendering of the overlay, red and green colors indicate bonds that are broken and created, respectively. The parts that are not taking part in the transformation are faded out in gray. Since the keto–enol tautomerization is an elementary reaction, the ITS is a cycle.⁵ See Supporting Information, Section A, for the case when formal charges are modified.

structure graphs before and after the chemical transformation, the ITS provides instructions for editing the reactant graph to obtain the product graph by means of a color-coded (edge symbols in the original ITS) representation of the difference between R and L. Obviously, the atom map must be known to edit a molecular structure or a pattern. The ITS, therefore, captures the same mechanistic information as a rule.

Obtaining rules is a challenge because the atom map is oftentimes not known, and reaction perception is notoriously difficult. Yet, reactions that are considered “elementary” are typically supported by mechanisms expressed as arrow-pushing diagrams, which directly yield the atom map. In this paper, we generate rules from the elementary steps of enzymatic mechanisms annotated with arrow-pushing diagrams and cataloged in the M-CSA database.⁶

2.1. Overlay Graphs. Consider a rule r_1 that is applied to reactants to generate products that become reactants of a rule r_2 , yielding further products. We can compress this reaction sequence into an “overall reaction” that collects all reactants and all products (including those generated by r_1 and ignored by r_2). The rule capturing this overall reaction is the composition of r_1 with r_2 .^{8,9} Rules r_1 and r_2 each have their own atom map, but to compose them, we need to track where in the products of r_1 there is a match for r_2 . Multiple matches may exist, and so r_1 and r_2 can give rise to several composite rules, reflecting distinct mechanisms.

Our objective here is to generalize the ITS idea from one elementary reaction to a sequence of elementary reactions. As an example, consider the enzymatic multi-step mechanism depicted in Figure 2. (Here the $-R_1$, $-R_2$, and $-R_3$ symbols stand for actual structures that we abbreviated for the sake of less clutter.) Each step is represented as an arrow-pushing diagram and can be translated into a rule. The first two steps are shown in the top row of Figure 3, where the middle pattern depicts the overlap between the right side of the rule for step 1 and the left side of the rule for step 2. The bottom row shows the corresponding composite rule. The intuition behind rule composition is perhaps best illustrated with the help of the ITS drawings associated with each elementary step, as shown in Figure 4A. The ITS of step 2 is overlaid (red dashed arrows) onto that of step 1. The overlay makes the ITS colors “interact” when modifications done by rule 1 are undone by rule 2. Such modifications are transient. For example, a bond is broken only to be created again, and charges that change in one step are restored in another. This is recorded by the color interaction. If only one ITS element is colored (no interaction), the color is simply inherited by the overlay: a bond or an atom is shown in green when the net change is a creation and in red when the net change is an elimination. If red overlays on green, the color changes to blue, indicating bonds that are created and then eliminated again. If green overlays on red, the color turns purple, indicating bonds and atoms whose original state is altered and then restored. “State” here refers to bond types (single, double, and triple) or charges. We refer to multiple ITS overlays as an OG. The OG corresponds to rule composition.

This process generalizes to n steps. While fairly intuitive, the proper definition and automation of rule composition require formalization, which we spell out in detail in the Supporting Information (Section A.1). The treatment presented there is meant to be a formal specification to guide implementation. At the time of this writing, our graph-transformation platform for chemistry⁴ does not handle the non-bond constraints (dotted lines) in full generality. The procedure with which we obtained the results in this paper is therefore a custom-implementation (script available online¹⁰) that makes use of the extant platform but does not extend it. A fully general treatment of chemical graph transformation as proposed by Behr et al.¹¹ is a task for the future.

The OG (or, equivalently, the composite rule) underscores the mechanistic significance of molecular parts that appear unchanged at the resolution of an overall reaction. The utility of a composite rule becomes apparent upon considering that it can be applied to a mixture other than the original reactants. A match of the L-pattern of the composite rule guarantees that the chemical transformation of the new reactants is compatible with the mechanism represented by the rule. Composite rules, therefore, enable the search for potential reactants other than those from which they were abstracted. In particular, they can suggest the assignment of a mechanism to a reaction for which none is known. This differs from a previous approach¹² that searched for a combination of elementary steps to achieve the desired outcome. While the previous approach can propose novel mechanisms by recombining steps, it suffers from the usual combinatorial explosion, confining it to rather short mechanisms. Here, the composite rule is a unit whose underlying steps can no longer be shuffled, thus not only curtailing the combinatorial explosion when searching for

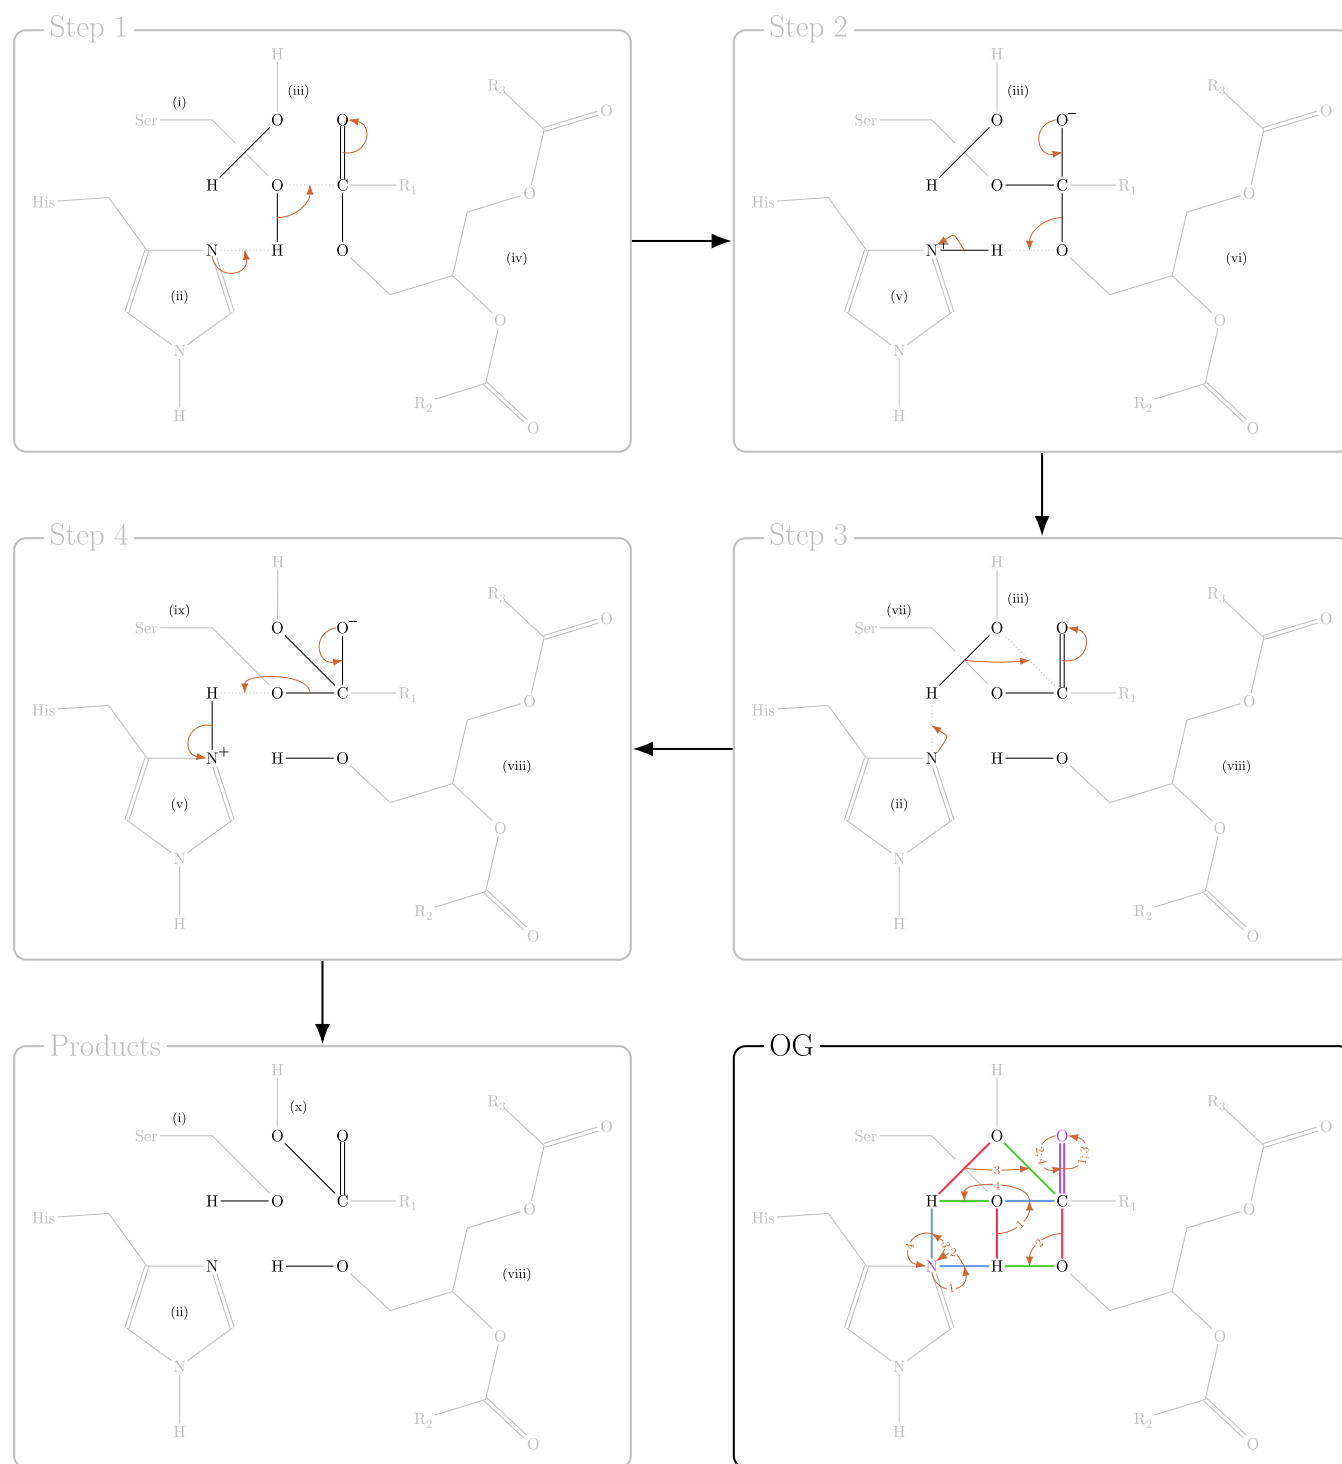


Figure 2. Four-step mechanism (M-CSA entry 218) of pancreatic triacylglycerol lipase (UniProtKB: P29183). This EC 3.1.1.3 reaction catalyzes the conversion of a triacylglycerol (iv) into a fatty acid ion (x) and diacylglycerol (viii). In the first step, a histidine (ii) deprotonates a serine (i), which now attacks a carbonyl carbon of the substrate (iv). As a result, the serine becomes covalently bound to the substrate (vi), and the histidine side chain is protonated (v). In the second step, the oxyanion collapses, releasing diacylglycerol (viii) and deprotonating the histidine (ii). In the third step, the histidine (ii) deprotonates a water molecule (iii), enabling the water to attack the acylated serine (vii). In step four, the oxyanion collapses once more, ejecting a fatty acid ion (x) and reconstituting the chemical state of the amino acid side chains. The lower right panel depicts all the steps stacked on top of each other, corresponding to the OG, as explained in the main text. Bonds created or broken during the reaction are shown in green and red, respectively. Transient bonds that are first created and then broken are shown in blue, whereas transient bonds that are first broken (or modified) and then restored are shown in purple. Likewise, atoms that undergo transient state changes (such as partial charge) are also shown in purple. Atoms and bonds in gray are not directly involved in the reaction and are shown for context.

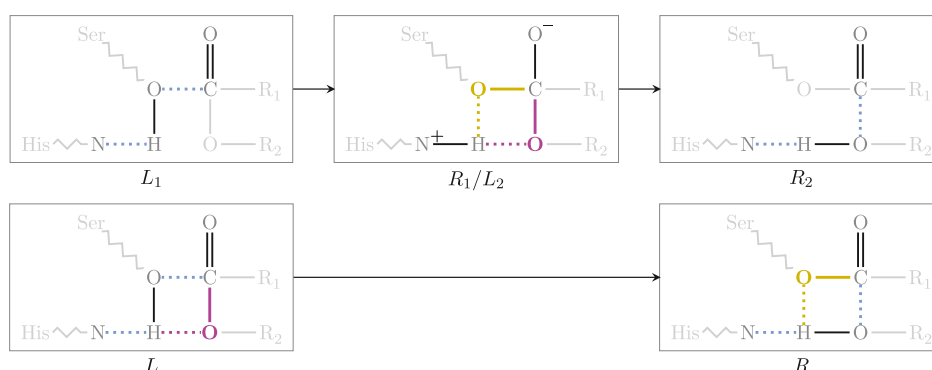


Figure 3. Rules $p_1 = L_1 \rightarrow R_1$ and $p_2 = L_2 \rightarrow R_2$ corresponding to steps 1 and 2 of the mechanism in Figure 2. The overlap of R_1 and L_2 is given by the relative position of the arrow pushes involved in the two steps. Yellow and magenta are used to highlight the sub-patterns unique to p_1 and p_2 , respectively. The gray parts are not part of the rules and are only included to situate them in the mixture. The bottom row shows the rule obtained by the composition of p_1 and p_2 . Conceptually, the left pattern L of the composite rule p must collect the conditions (the structural parts) required by the left pattern of p_1 plus those conditions of the left pattern of p_2 that have not been created by the action of p_1 . Likewise, the right pattern R of p must contain the right pattern of p_2 plus those parts of the right pattern of p_1 that have not been modified by p_2 . This is formalized in Supporting Information Section A.1.




Figure 4. (A): Graphical superposition of ITS drawings associated with the rules derived from steps 1 and 2 of the mechanism in Figure 2. The color coding is explained in the main text. The gray parts do not play a part in the superposition and are only used to localize the ITS elements within the reactants. (B) Resulting OG obtained by superimposing steps 1 and 2.




Figure 5. (A) OG of the triglyceride lipase mechanism from Figure 2. The color coding is detailed in the text. The gray parts of the molecules are only shown for context and are not components of the OG or the associated composite rule. In the present example, the catalyst is present only through a nitrogen atom of the histidine and the OH group of the serine. (B) sOG derived by removing all atoms belonging to the catalyst. See the text for details.

compatible reactant/product pairs but also limiting the present approach to known mechanisms. We illustrate applications in the section “Results and Discussion”.

A few remarks are in order to keep perspective. Our framework represents a microscopic view in which atoms have unique identifiers, which is distinct from a more coarse-

grained view based on chemical equivalence. This affects the meaning of “change” or “modification”. For example, a bond remains (overall) unchanged only if it is broken and reconstituted between the same partners at the microscopic level. This is useful in the context of enzyme reactions because it detects atoms exchanged between the enzyme and




Figure 6. (A) Composite rule of the full triglyceride lipase mechanism from Figure 2. The catalytic components of the composite rule are highlighted in yellow. Notice that the atoms are not required to preserve their catalytic identity across the reaction. (B) Corresponding substrate rule obtained by removing the highlighted catalytic components. Red and green are used for parts that are removed and, respectively, created by the rule. The atoms that change their catalytic identity appear as being destroyed or created in the substrate rule.

its substrates, although the enzyme chemically remains the same.

The OG is guaranteed to capture all atoms and bonds transiently involved in the reaction mechanism. Yet, the present formalism does not retain the exact nature of transient modification in all cases. For instance, in Figure 5A, the double bond is flagged as transiently modified, but the more specific information that it has been reduced to a single bond twice is lost. Similarly, the N of the histidine transiently changed charge, but the number of times it changed and the direction of state change (whether charge increased or decreased) are not recorded. In the final analysis, from a composite rule, as formalized here, we cannot recover the following information: (i) the history of temporary changes a bond underwent before restoration of the original state; (ii) the history of transient changes a bond underwent before a non-transient change; (iii) the transient types (single, double, and triple) of a bond formed and eliminated again (blue bond in the OG). We detail these limitations in Supporting Information, Section D.

2.2. Substrate Rules. In the case of catalysis, we can make a useful generalization by abstracting the identity of the catalyst in the composite rule to only keep those molecular parts of the substrate that are required in the interaction with the catalyst, given the mechanism captured by the rule. In other words, we can “anonymize” the catalyst without compromising the mechanism. By way of metaphor, consider a clip of a dance movement that requires two parties, say a waltz, but make one of the partners transparent to show only the movements of the other. In such a “minus one” clip, any other partner making waltz movements can be dropped in place of the transparent one.

More specifically, following Figure 2, consider the hydrogen atom that belongs to the serine of the catalyst. This H is given to the substrate, and a different H, originally belonging to a water molecule and transiently parked at the N of a histidine, is returned to serine, restoring its chemical state. Clearly, this H is essential for the reaction, but it does not strike one as essential that it must come from an O. Any proton donor would do. A temporary parking space for hydrogens is essential for this mechanism, but, again, it seems less important that the parking lot be a N. The broader context—that the O belongs specifically to a serine and the N to a histidine—is not represented in any way in these rules, at least not at the present level of abstraction. Likewise, the proton donor and the temporary parking lot need the right pK_a to engage in catalysis, but this, too, is not within the scope of the simple graphical model used here.

These observations suggest a reduced (thus more general) composite rule in which only those parts of the substrates are retained that are required for their transformation. The enzyme and any other catalytic molecules are abstracted away. We refer to this reduced rule as a substrate rule or, in the OG representation, as the substrate OG (sOG). The substrate rule has the advantage of being applicable to a mixture consisting only of substrate molecules, leaving the exact nature of the catalyst(s) unspecified yet retaining the structural requirements from the mechanism. This means that any alternative catalysts that can be plugged into the reaction (i.e., match the substrate rule) are potential catalysts for the transformation of the same substrate using the same mechanism.

To construct the substrate rule associated with a catalytic mechanism, we need to remove from the composite rule those parts representing the catalyst. This is illustrated in Figure 6A for the whole mechanism shown in Figure 2 (This composite rule differs from the one shown in Figure 3, which resulted from composing only step 1 and step 2.).

We follow the usual notion of a catalyst as a molecule that appears both on the educt and product sides. The molecular parts of the catalyst have to be identified separately on the educt and the product side. These parts are highlighted in gold in Figure 6A. For instance, the hydrogen atom that is part of the serine before the reaction is different (in terms of microscopic identity) from the one attached afterward. The substrate rule obtained by removing the highlighted atoms and bonds from the composite rule is shown in Figure 6B, and the corresponding sOG is shown in Figure 5B.

As a consequence of this resection, the left and right patterns of a substrate rule no longer have to contain the (microscopically) same atoms. Yet, mass conservation remains intact because *L* and *R* of the substrate rule contain the same number and type of atoms. For example, the substrate rule of Figure 6B means that the H on the left, colored red, literally vanishes from the substrate system, and the H on the right, colored green, appears out of nothing. This is a consequence of having removed the catalyst from the picture. In a real-world setting, the H that disappears must be received by a catalyst, while the H that seems to be created out of nothing must be provided by the same catalyst. The point of the substrate rule is to express that the H donor and acceptor can belong to any catalyst, not just the particular one whose mechanism gave rise to the composite rule. Moreover, if the H-donor and H-acceptor were not catalysts, then some additional molecules would have been modified and would have to show up in the rule. The

substrate rule thus guarantees the catalytic nature of transformation without specifying the catalyst.

We described the “typical” case and, for the sake of readability, glossed over some fine print, on which we elaborate in [Supporting Information](#), Sections A.2 and E. Some of that fine print is purely technical, but some is conceptually interesting and worth a quick brush stroke. The formalization of the substrate rule tugs at the question of what, exactly, qualifies as a catalyst. Consider a situation in which an enzymatic process generates an intermediate product X that then serves in a catalytic capacity for the remainder of the process by, say, helping restore the state of enzyme E. The overall reaction might read like $E + S \rightarrow E + P + X$. Clearly, X is not a catalyst by IUPAC standards since it does not show up on the educt side. Yet, in the context of the mechanism, X participates as a catalyst generated in situ. The substrate rule formalism is not guaranteed to work in such cases without additional assumptions. An example, where X is H_2O , is given in Section E of the [Supporting Information](#). (With a bit of squinting, Section E can be read without plowing through the formalism.)

3. RESULTS AND DISCUSSION

3.1. Obtaining OGs. The M-CSA database⁶ is a non-redundant catalogue of enzymatic reaction mechanisms represented as sequences of elementary steps annotated with electronic displacements in the style of arrow pushing. This kind of information is suited for constructing OGs, as in [Figure 2](#).

One problem is that the reaction steps listed in the M-CSA do not explicitly specify atom maps. This does not present an issue for each step considered in isolation since the flow of electronic displacements suffices to relate atoms in the educts to those in the products. Rather, the problem is the absence of atom maps across steps. While the visual display of an M-CSA mechanism might suffice for a human user to infer the connections between steps, visualization comes with connotations that are not always present in the underlying data and therefore not available for automated analysis.

The absence of a specific atom map connecting the steps of an M-CSA mechanism opens the possibility for the left pattern of one step to match the products of the previous step in several ways, should these products contain multiple equivalent instances of a chemical substructure, such as the hydrogen atoms of a methyl or an amino group. Likewise, the presence of several instances of the same molecular species, like H_2O , during a mechanism yields a choice of which instance to use in a given step. Such symmetries are relevant for the construction of the composite rule only if they involve substructures with distinct histories. For example, [Figure 7](#) illustrates how the atom map linking two steps impacts the composite rule as represented by its OG. The first step attaches a hydrogen to an amine group, H_2N , while the second cleaves a hydrogen atom from the resulting H_3N^+ group. Since all three hydrogen atoms of that group are equivalent, the left pattern of the second rule can match any one of them. However, since the hydrogen atoms that came with the initial H_2N have necessarily acquired a different history from the hydrogen atom that was just added to form H_3N^+ , we obtain two non-isomorphic mechanisms and associated OGs, depending on which class the cleaved hydrogen belongs to.




Figure 7. Left pattern of a rule removing a hydrogen atom can be compatible with any of the three hydrogen atoms of the H_3N^+ group. Depending on the choice of match, two distinct rule compositions are possible, resulting in different OGs.

Rule composition is based on an underlying sequence of rule applications (the mechanism), each of which involves a particular choice of match into the mixture. A specific sequence of matches, along with the atom identities across each step, defines the atom map for the entire mechanism. To avoid confusion, we refer to the series of steps in an M-CSA entry as a “sequence” and reserve the term “mechanism” for a sequence with an atom map. Since, for many M-CSA entries, several matches are possible at each step, we consider all of them and treat them as distinct mechanisms.

At the time of writing, the M-CSA database contains 888 sequences of steps.¹³ We were able to generate 2802 mechanisms, with attendant composite rules and OGs, for a total of 600 M-CSA sequences. 288 M-CSA sequences could not be cast as rules due in part to current limitations of the implementation of chemical graph transformation, which, for example, does not yet process radical formation, and in part due to difficulties with extracting the needed information from some M-CSA entries. For the process used to prepare the M-CSA information for our approach, see.¹² Of these 600 M-CSA sequences, 435 have a unique mechanism and 100 have two mechanisms. Five M-CSA sequences are compatible with over 100 different mechanisms. One sequence, M-CSA entry 38, gives rise to 964 non-isomorphic composite rules because it involves four water molecules that can become distinguishable as the mechanism proceeds.

The number of possible non-isomorphic rule compositions grows exponentially with the number of disjoint isomorphic substructures. Thus, even a partial atom map that breaks some symmetries can reduce the number of OGs substantially. For example, M-CSA entry 186 has 24 possible mechanisms that can be reduced to a single OG by mapping just four atoms, each in one intermediate.

3.2. Interpretation of OGs. The construction of a composite rule, and by implication its OG, preserves information about transient changes. If a mechanism includes some steps that are independent of others, their relative order of temporal execution is not fixed. Naturally, the compression of a sequence of steps into a single overall rule loses information about the ordering of such steps, but by their independence, this ordering is irrelevant for the outcome. Independence is often the case for “clean-up” steps that reconstitute the chemical state of the enzyme. At the end of [Section 2.1](#), we summarized the kind of information that is lost by the composite rule or the OG.

To determine whether the information preserved in an OG provides insights into enzymatic reaction mechanisms, we constructed the OGs of reactions within the EC 3.1.1.-

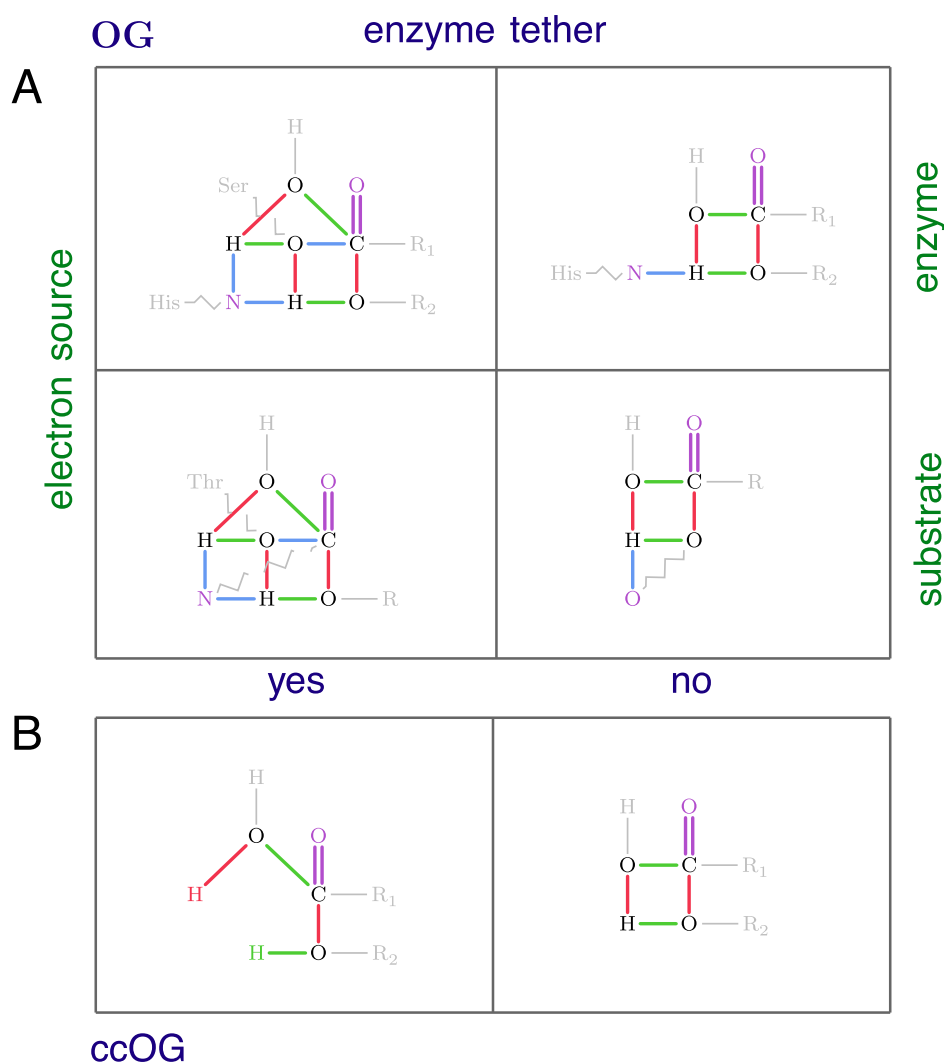


Figure 8. (A) Classification of different OG patterns extracted from EC 3.1.1.- reactions listed in the M-CSA. (B) Catalytic complement OGs derived from OGs by removing catalytic sites regardless of whether they belong to a distinct catalyst molecule or the substrate or both. These ccOGs distinguish between the equivalence classes along the tether/no-tether axis.

subclass, of which the mechanism depicted in Figure 2 is an example. The first position in the EC numbering scheme (i.e., EC 3.-.-) refers to the hydrolytic cleavage of bonds. This process transfers a substrate group to specifically H₂O as the acceptor. The second position (EC 3.1.-) indicates the nature of the hydrolyzed bond: in the present case, an ester bond. The third position (EC 3.1.1.-) specifies the nature of the substrate, here a carboxylic ester. The M-CSA lists 24 reaction sequences as EC 3.1.1.-. Converting these sequences into mechanisms does not reveal much ambiguity: only two reaction sequences allow for more than one OG—specifically, 2 and 5 OGs—yielding a total of 29 OGs. The ambiguities arise from the equivalence of hydrogen atoms in an H₂O⁺ group and an H₃N⁺ group, similar to the case illustrated in Figure 7.

The set of EC 3.1.1.- reactions indeed translates into patterns at the level of OGs (Figure 8), revealing two mechanistic distinctions that organize enzyme–substrate relations into four classes.

The first distinction refers to how many catalytic sites engage in a chemical interaction with the substrate. In the context of EC 3.1.1.- reactions, this amounts to whether two

or one hydrogen atom are involved in the catalytic process, which in turn indicates whether or not the mechanism proceeds via a covalently bound intermediate. In the case of a bound intermediate, the flow of electronic displacements repeats twice: once to create a covalent bond with the enzyme, followed by the release of one product, and a second time to break that covalent bond, replacing it with a bond to the OH of water, thereby releasing the second product of hydrolysis. Each flow is a round trip of electrons from a source to a sink and back, each round trip involving a distinct proton that engages in a transient bond. Absent a covalently bound intermediate, only one round trip is necessary, requiring only one proton.

The second distinction pertains to whether the source of electrons that initiates the reaction sequence is located on the enzyme or the substrate. At the abstraction level of arrow pushing, an elementary step is a concerted electronic displacement prompted by a region of high electron density sufficiently “motivated” by a chemical environment. Thus, a special role falls to any step in a catalytic process that initiates an electronic displacement sequence from a relatively stable

intermediate, including, in particular, the very first step of the reaction.

Figure 8 arrays the four OG patterns according to electron source and whether tethering of an intermediate occurs. The upper left quadrant is our running example of Figure 2, entry 218. The reaction mechanism generates a covalently bound intermediate with the enzyme as the initial electron source. This pattern is present in 19 of the 29 EC 3.1.1.- OGs. These 19 OGs are derived from as many M-CSA reaction sequences with no ambiguity in the underlying mechanisms. The lower left quadrant shows the case of a D-aminoacyl-tRNA deacylase, entry 748 (sequence proposal #1), in which the tethering occurs to a threonine of the enzyme, but the electron source is provided by the N of an amino group bound to a C in α position of the carbonyl group of the substrate. The upper right quadrant shows entry 83, phospholipase A2, whose hydrolase action proceeds without creating a covalently tethered intermediate and uses a histidine residue as the initial electron source. Although no covalent tethering occurs, the substrate is nonetheless locked in a position by coordination to Ca^{2+} . At the abstraction level of OGs, such bonds are not considered as properly catalytic with regard to the actual chemical transformation but rather as screening and/or securing a substrate in the catalytic pocket. These interactions are presumably critical for catalytic efficiency and thus belong to the catalytic process in a broader sense that could, in principle, be captured by rules and attendant OGs with substantially more context beyond the reaction center. Yet, such extensive and specific rule refinements would blur the general picture, making OGs less useful. Finally, the lower right quadrant is again the case of the D-aminoacyl-tRNA deacylase, entry 748 but this time based on an alternative reaction sequence proposal listed in the M-CSA. In this case, the hydrolysis is entirely intramolecular to the substrate, with the enzyme just providing a tailored physico-chemical environment. This particular reaction sequence gives rise to two mechanisms and thus OGs due to the equivalent hydrogen atoms in the protonated OH-group after the first step. It should be noted, however, that the M-CSA lists the reaction sequences for entry 748 with a low confidence rating.

Especially in the context of enzymatic reactions, the notion of a catalyst typically connotes an entity distinct from the substrate. The above observations, however, show that catalytic functionality can be distributed across molecular entities and may, therefore, cross-cut the traditional separation of the enzymes and substrates, even though known cases seem rare so far. If we drop the notion of a unitary catalyst in favor of a collection of catalytic sites that could be distributed across the theater of a chemical reaction, then our construction of removing atoms carrying catalytic functionality from the OG to produce a sOG would, in the case of the lower quadrants in Figure 8A, remove atoms belonging to the substrate itself. This generalization suggests calling the resulting construct a “catalytic complement” OG or ccOG (Figure 8B) rather than a sOG. The ccOG of the upper left quadrant is the same as the sOG since the substrate carries no catalytic sites. Yet, the ccOG of the lower left quadrant is isomorphic to that of the upper left quadrant. Similar observations hold for the two right quadrants. The ccOG appears to classify mechanisms according to the number of catalytic sites involved.

It is tempting to speculate that the evolution of some enzymes might involve shifting catalytic functionality from substrates to enzymes or vice versa. For example, the efficiency and reliability of an originally self-catalyzed intramolecular reaction (such as the lower right quadrant) might be improved by moving sites of catalysis from the substrate to a polypeptide pocket that originally provided no more than a suitable environment for the reaction.

3.3. Using Substrate Rules to Identify Compatible Reactions. The applicability of a composite rule to a mixture guarantees that the chemical mechanism represented by the rule is feasible in the mixture. Feasibility must be understood at the abstraction level defined by graph transformation, which, as it stands, has nothing to say about thermodynamics. With this in mind, we can use a composite rule to screen for reactions whose educts and associated products match the rule's left and right patterns, respectively, to suggest a mechanism when none is known or propose an alternative to a known mechanism. To be clear, a composite rule is applicable much like any other rule to a variety of reactants matching its left pattern. The point here is that the composite rule represents a series of steps and can be constructed programmatically, as shown in earlier sections, thus automating the search for educt–product pairs that are compatible with a given mechanism.

To this end, we extracted over 13,000 overall reaction descriptions from the Rhea database, an expert-curated database of biologically interesting reactions.⁷ Since these descriptions do not explicitly list catalytic components, we deploy substrate rules, which capture the mechanism while minimizing the specification of the catalyst. To test the approach, we use substrate rules derived from the composite rules that represent the instances in the EC 3.1.1.- reaction class of the M-CSA. The 29 composite rules and associated OGs characterized in the previous section are reduced to 11 equivalence classes at the level of substrate rules and sOGs: one class of size 18, one of size 2, and nine of size 1. Of these 11 distinct substrate rules, 3 are refinements of others; that is, to say, in 3 cases, when one rule applies, another one, strictly less selective, applies as well and produces the same result. We therefore discard the more selective rules. One substrate rule came from *para*-nitrobenzyl esterase, which cleaves a C–N bond, making its EC 3.1.1.- classification uncertain. We opted to leave it out, ending up with 7 distinct substrate rules capturing EC 3.1.1.- mechanisms.

The M-CSA and Rhea databases use different conventions for carboxylic acids. For each reaction marked as EC 3.1.1.-, the M-CSA writes a carboxylic acid group as COOH, whereas Rhea writes it as COO^- plus H^+ . Since our rules are derived from the M-CSA, we need to additionally cleave the O–H bond of the carboxylic acid in the product. Recall that this OH group comes from a water molecule on the educt side. In rules that are minimal in the sense of only representing the reaction center, the water molecule need not be fully specified. For example, the rule that captures the mechanism of M-CSA entry 218 and is shown in Figure 6A specifies one H–O bond of water in its left pattern but not the other O–H bond, even though the latter ends up in a product molecule. The reason is that nothing happens to that other O–H bond; it simply gets transferred along as a result of the creation of a bond between the carbon of the carboxyl group and the O of the water. However, if that “inherited” bond is cleaved in a subsequent step in order to achieve the

carboxylate form, the modified rule would then have to fully specify the water molecule. Since a larger left pattern makes the rules more specific, we also keep the 7 original unmodified rules.

These 14 substrate rules match 588 Rhea reactions, which are, therefore, compatible with the catalytic mechanism encapsulated in the matching rules. We utilize these matches for two types of analyses: to evaluate the retrieval rate of the known EC 3.1.1.- reactions and to explore reactions of different or unknown EC classes which are compatible with an EC 3.1.1.- mechanism.

A total of 130 reactions are marked as EC 3.1.1.- within the Rhea database. Out of these, we retrieve 90, giving us a recall rate of 69.2%. All these matches are, as expected, against the modified rules. The reasons for not matching the remaining 40 reactions are manifold. 20 reactions exhibit an oligomeric nature not reflected in the SMILES representation from which we constructed the molecule graphs, or they show educts and products using mismatched abstractions or tautomerizations; 5 reactions involve interactions with iron, which are not supported by our present version of the graph-transformation approach to chemistry. The remaining 15 reactions do not stop after cleavage of the carboxylic ester bond but exhibit “follow-up” reactions of the kind listed in Table 1.

Table 1. Types of Additional Modifications to the Product Molecules of the Carboxylic Ester Bond Cleavage Observed in the Rhea Database with Incidence

description	incidence
keto–enol tautomerization	5
second carboxylic ester bond cleavage	3
decarboxylation	3
cyclization	2
intra-molecular proton transfer	1
release of the sulfite group	1

Among these follow-up reactions, we sought to include the keto–enol tautomerization as the most prevalent. The M-CSA database contains two elementary reaction steps corresponding to keto–enol tautomerization, step 4 in entry 85 and step 4 in entry 463, which are both described as spontaneous tautomerization occurring outside the active site of the enzyme. We therefore augmented our rules by composing them with the steps corresponding to spontaneous keto–enol tautomerization as present in the M-CSA. These steps were assumed to involve the ester oxygen, which ends up as a ketone group instead of an alcohol group. The augmented rules now match four of the five Rhea reactions in this category. In the remaining reaction, the keto–enol tautomerization occurs twice; one involves the former ester oxygen and the other involves an alcohol group already present in an educt molecule. Other than that, the augmented rules do not match more Rhea reactions.

From among the remaining 498 matched reactions, not marked as EC 3.1.1.-, 30 are given a different EC number in the Rhea database and 468 reactions are assigned no EC number in Rhea. These results are summarized in Table 2.

The majority of reactions classified as different classes than EC 3.1.1.- are transferases, which cleave an ester bond and create a new one using an alcohol group of another educt molecule, as in $R-CO-OR' + OH-R'' \rightarrow R-CO-OR'' +$

Table 2. EC Classification, as Declared in the Rhea Database, of the Reactions Explained by Our Substrate Overlay Rules

EC class	incidence	description
none	468	
EC 3.1.1.-	90	carboxylic ester hydrolases
EC 2.3.1.-	19	acyltransferases other than aminoacyl
EC 3.6.1.-	3	phosphorus-containing acid anhydride hydrolases
EC 2.3.2.-	2	aminoacyltransferases
EC 4.2.99.-	2	other carbon–oxygen lyases
EC 5.4.1.-	2	intra-molecular acyl transferases
EC 2.1.3.-	1	carboxy-/carbonyltransferases
EC 3.2.1.-	1	glycosidases

$OH-R'$. The alcohol group acts in place of the OH of water in the hydrolase context, where $R'' = H$. These reactions resulted from a match to the unmodified substrate rules since the water molecule is fully specified in the modified version.

Cross-linking the 468 reactions of the unspecified EC class with UniProt¹⁴ reveals 310 reactions associated with at least one enzyme classified as EC 3.1.1.-. 200 of those are exclusively associated with EC 3.1.1.- enzymes, and 26 reactions are associated with no EC 3.1.1.- enzyme in UniProt. Similar to the results in Table 2, most of the retrieved enzymes outside of EC 3.1.1.- are acyltransferases (EC 2.3.1.-). Other prominent classes contain thiolester hydrolases (EC 3.1.2.-) and carbon–nitrogen hydrolases (EC 3.5.1.-). We retrieved no UniProt reference for the remaining 132 reactions.

4. SUMMARY AND CONCLUSIONS

Our starting point is the representation of chemistry within the framework of graph transformation, enabled by an open software platform developed for this purpose.⁴ The widespread availability of versatile open-source tools for organic chemistry (e.g., RDKit¹⁵) based on a linear string encoding of molecular structures begs the question of why a framework based on graph transformation is needed. An answer implicit in the present contribution is that the graph-transformation approach forces the inclusion of the atom map. By virtue of the atom map, rules (of graph transformation) capture a mechanistic aspect of chemical reactions, albeit a very limited one. Nothing is said about the ordering of electronic displacements for complex mechanisms or about quantitative aspects that depend on geometry and thermodynamics. Because mathematical graphs are not two-dimensional structures, they simply represent connectivity. While the framework can, in principle, be augmented to include information about quantitative aspects, the current approach has already demonstrated theoretical and practical utility,^{12,16–21} to which the present work seeks to add.

We construct the composite rule of a sequence of rules derived from elementary steps—elementary in the sense of arrow pushing. The composite rule represents the “overall reaction”. A composite rule records not only the net changes between educts and products but also transient changes, such as the temporary making or breaking of a bond or the temporary loss or acquisition of a formal charge. We visualize the information retained by a composite rule in terms of a single graph—the OG—that superposes educt and product patterns along with a color-coded “fate map” of bonds reconstructed from the composite atom map. The OG can be

viewed as the representation of a reaction mechanism using a graph-edit notation and constitutes a generalization to multi-step reactions of the ITS proposed by Fujita for elementary reactions.⁵

We pursue two specific applications in the context of enzymatically catalyzed reactions, enabled by a valuable resource of enzyme mechanisms known as the M-CSA. In the first case, we automatically construct the OGs of all carboxylic ester hydrolase mechanisms (EC 3.1.1.-) listed in the M-CSA that have sufficient information to be translated into our framework. The compact nature of OGs enables visual inspection, which made us realize that these graphs can be classified according to whether the mechanism proceeds by tethering an intermediate product and whether the electronic displacement step initiating the reaction sequence resides within the substrate or the enzyme. This suggests that OGs could be a useful and interpretable data type for processing complex reactions with statistical learning algorithms.

Any reaction whose educts and products match a composite rule can proceed by the mechanism the rule encapsulates. This provides a means for suggesting mechanisms of catalysis for reactions whose mechanisms are unknown or for whom alternative mechanisms are sought. Since the composite rule includes catalytic parts that interact with the reaction center, candidate reactions listed in databases without specification of the catalyst would not be matched. To this end, we generalize the composite rule by eliminating all references to the catalyst while retaining the “footprint” of the catalytic action on the portion of the reaction center belonging to the substrate. With this construct, termed substrate rule (and the associated sOG), we can identify reactions with unknown mechanisms or unknown catalysts that could proceed with the captured mechanism. We exemplify this use case by screening all of the reactions in the Rhea database using the substrate rules derived from EC 3.1.1.- reactions in the M-CSA collection. We recovered a substantial amount of reactions that were indeed classified as EC 3.1.1.- in Rhea, alongside many reactions that had no classification but whose study suggested compatibility with an EC 3.1.1.- mechanism. This illustrates the power of rule composition for automating the search for reactions that qualify as candidates for a given mechanism.

To end on a conceptual note, the idea behind the substrate rule is to capture the requirements of a specific catalyzed mechanism while being as agnostic as possible about the chemical implementation of the catalyst. This often works, but the graph-transformation framework requires deciding who the catalyst is. A catalyst is usually viewed as a unitary object present on the educt side of a reaction and reconstituted on the product side. If all cases were like this, the construction of the substrate rule would pose no problem. A difficulty arises, however, in cases in which “the” catalyst is actually a collection of catalytic sites that are distributed across the reaction theater, including the substrate itself. Even more vexing is the case of multi-step reactions, where, in addition to a catalyst present initially, a molecule (or functional group) is formed at an intermediate step of the reaction process and subsequently participates as a catalyst in completing the reaction. Such a “half-time” catalyst would show up as a product of the overall reaction but not as an educt and thus not qualify as a catalyst despite its crucial catalytic contribution.

DATA AND SOFTWARE AVAILABILITY

All utilized data, in the form of arrow-pushing diagrams and molecule smiles, comes from publicly available databases, M-CSA⁶ and Rhea.⁷ The code for converting the reaction sequences in M-CSA into graph-transformation rules compatible with MØD⁴ is available here: https://github.com/chrstf/mcsa_rule_converter. The code for computing the OGs and substrate rules from M-CSA data is available here: <https://github.com/JuriKolcak/overlay-graphs>. Searching for reactions compatible with a given composite rule has been conducted using a version of MØD specifically augmented for better performance with rules featuring multiple unconnected components,²² available here: <https://github.com/jakobandersen/mod/tree/archive/rule-application-21>.

ASSOCIATED CONTENT

Supporting Information

The Supporting Information is available free of charge at <https://pubs.acs.org/doi/10.1021/acs.jcim.2c00426>.

A thorough introduction of the graph-transformation formalism with chemical examples; category theoretical formalism of rule composition and substrate rule construction; example of a non-hydrogen atom exchanged between the substrate and the catalyst; limits of the rule composition approach to OGs; and discussion on the nature of catalysts in light of substrate rules (PDF)

AUTHOR INFORMATION

Corresponding Authors

Juri Kolčák – Department of Mathematics and Computer Science, University of Southern Denmark, 5230 Odense, Denmark; Department of Systems Biology, Harvard Medical School, Boston, Massachusetts 02115, United States; orcid.org/0000-0002-9407-9682; Email: juri.kolcak@gmail.com

Daniel Merkle – Department of Mathematics and Computer Science, University of Southern Denmark, 5230 Odense, Denmark; orcid.org/0000-0001-7792-375X; Email: daniel@imada.sdu.dk

Authors

Jakob L. Andersen – Department of Mathematics and Computer Science, University of Southern Denmark, 5230 Odense, Denmark

Rolf Fagerberg – Department of Mathematics and Computer Science, University of Southern Denmark, 5230 Odense, Denmark; orcid.org/0000-0003-1004-3314

Christoph Flamm – Department of Theoretical Chemistry, University of Vienna, 1090 Vienna, Austria

Walter Fontana – Department of Systems Biology, Harvard Medical School, Boston, Massachusetts 02115, United States; orcid.org/0000-0003-4062-9957

Christophe V.F.P. Laurent – Department of Mathematics and Computer Science, University of Southern Denmark, 5230 Odense, Denmark; orcid.org/0000-0002-9112-6981

Nikolai Nøjgaard – Department of Mathematics and Computer Science, University of Southern Denmark, 5230 Odense, Denmark

Complete contact information is available at:

<https://pubs.acs.org/10.1021/acs.jcim.2c00426>

Notes

The authors declare no competing financial interest.

ACKNOWLEDGMENTS

All authors are supported by the Novo Nordisk Foundation grant NNF19OC0057834. J.L.A., R.F., C.F., and D.M. are additionally supported by the Novo Nordisk Foundation grant NNF21OC0066551. J.L.A., R.F., W.F., and D.M. are further supported by the Independent Research Fund Denmark, Natural Sciences, grant DFF-0135-00420B. Finally, J.L.A., R.F., and D.M. are supported by the Independent Research Fund Denmark, Natural Sciences, grant DFF-7014-00041.

REFERENCES

- (1) Levy, D. E. *Arrow-Pushing in Organic Chemistry*; John Wiley & Sons, 2011; p 320.
- (2) Ehrig, H.; Ehrig, K.; Golas, U.; Taentzer, G. *Monographs in Theoretical Computer Science. An EATCS Series*; Brauer, W., Rozenberg, G., Salomaa, A., Eds.; Springer: Berlin and Heidelberg, Germany, 2006.
- (3) Habel, A.; Müller, J.; Plump, D. Double-pushout Graph Transformation Revisited. *Math. Struct. Comput. Sci.* **2001**, *11*, 637–688.
- (4) Andersen, J. L.; Flamm, C.; Merkle, D.; Stadler, P. F. A Software Package for Chemically Inspired Graph Transformation. *Graph Transformation. ICGT 2016. Lecture Notes in Computer Science*; Echahed, R., Minas, M., Eds.; Springer: Cham, Switzerland, 2016; Vol. 9761; pp 73–88.
- (5) Fujita, S. Description of Organic Reactions Based On Imaginary Transition Structures. 1. Introduction of New Concepts. *J. Chem. Inf. Comput. Sci.* **1986**, *26*, 205–212.
- (6) Ribeiro, A. J. M.; Holliday, G. L.; Furnham, N.; Tyzack, J. D.; Ferris, K.; Thornton, J. M. Mechanism and Catalytic Site Atlas (M-CSA): a Database of Enzyme Reaction Mechanisms and Active Sites. *Nucleic Acids Res.* **2017**, *46*, D618–D623.
- (7) Lombardot, T.; Morgat, A.; Axelsen, K. B.; Aimò, L.; Hyka-Nouspikel, N.; Niknejad, A.; Ignatchenko, A.; Xenarios, I.; Coudert, E.; Redaschi, N.; Bridge, A. Updates in Rhea: SPARQLing Biochemical Reaction Data. *Nucleic Acids Res.* **2019**, *47*, D596–D600.
- (8) Behr, N.; Sobocinski, P. Rule Algebras for Adhesive Categories. *27th EACSL Annual Conference on Computer Science Logic (CSL 2018)*; Dagstuhl: Germany, 2018; pp 11:1–11:21.
- (9) Andersen, J. L.; Flamm, C.; Merkle, D.; Stadler, P. F. Rule Composition in Graph Transformation Models of Chemical Reactions. *MATCH Commun. Math. Comput. Chem.* **2018**, *80*, 661–704.
- (10) Kolčák, J.; Laurent, C. V. F. P.; Nøjgaard, N. Overlay Graph Source code. 2022, <https://github.com/Jurikolcak/overlay-graphs> (accessed on April 12, 2022).
- (11) Behr, N.; Krivine, J.; Andersen, J. L.; Merkle, D. Rewriting Theory for the Life Sciences: A Unifying Theory of CTMC Semantics. *Theor. Comput. Sci.* **2021**, *884*, 68–115.
- (12) Andersen, J. L.; Fagerberg, R.; Flamm, C.; Fontana, W.; Kolčák, J.; Laurent, C. V. F. P.; Merkle, D.; Nøjgaard, N. Graph Transformation for Enzymatic Mechanisms. *Bioinformatics* **2021**, *37*, i392–i400.
- (13) EMBL-EBI, M-CSA API. <https://www.ebi.ac.uk/thornton-srv/m-csa/api/entries/?format=json> (accessed Sept 24, 2021).
- (14) The UniProt Consortium. UniProt: the Universal Protein Knowledgebase in 2021. *Nucleic Acids Res.* **2020**, *49*, D480–D489.
- (15) Landrum, G. Open-source cheminformatics. <https://www.rdkit.org> (accessed Jan 21, 2022).
- (16) Silva, W. M. C. d.; Andersen, J. L.; Holanda, M. T.; Walter, M. E. M. T.; Brigido, M. M.; Stadler, P. F.; Flamm, C. Exploring Plant Sesquiterpene Diversity by Generating Chemical Networks. *Processes* **2019**, *7*, 240.
- (17) Andersen, J. L.; Flamm, C.; Merkle, D.; Stadler, P. F. In silico Support for Eschenmoser's Glyoxylate Scenario. *Isr. J. Chem.* **2015**, *55*, 919–933.
- (18) Andersen, J. L.; Merkle, D.; Rasmussen, P. S. Combining Graph Transformations and Semigroups for Isotopic Labeling Design. *J. Comput. Biol.* **2020**, *27*, 269–287.
- (19) Andersen, J. L.; Fagerberg, R.; Flamm, C.; Kianian, R.; Merkle, D.; Stadler, P. F. Towards Mechanistic Prediction of Mass Spectra Using Graph Transformation. *MATCH Commun. Math. Comput. Chem.* **2018**, *80*, 705–731.
- (20) Andersen, J. L.; Andersen, T.; Flamm, C.; Hanczyc, M. M.; Merkle, D.; Stadler, P. F. Navigating the Chemical Space of HCN Polymerization and Hydrolysis: Guiding Graph Grammars by Mass Spectrometry Data. *Entropy* **2013**, *15*, 4066–4083.
- (21) Andersen, J. L.; Flamm, C.; Merkle, D.; Stadler, P. F. Integer Transformation Motifs — Modelling Pathways as Integer Hyperflows. *IEEE/ACM Trans. Comput. Biol. Bioinf.* **2019**, *16*, 510–523.
- (22) Andersen, J. L.; Fagerberg, R.; Kolčák, J.; Laurent, C. V. F. P.; Merkle, D.; Nøjgaard, N. Efficient Modular Graph Transformation Rule Application. **2022**, arXiv preprint arXiv:2201.04360.

Recommended by ACS

A Bayesian Approach to Extracting Kinetic Information from Artificial Enzymatic Networks

Mathieu G. Baltussen, Wilhelm T. S. Huck, *et al.*

MAY 12, 2022
ANALYTICAL CHEMISTRY

READ 

Quantum Mechanics and Machine Learning Synergies: Graph Attention Neural Networks to Predict Chemical Reactivity

Mohammadamin Tavakoli, Pierre Baldi, *et al.*

JANUARY 12, 2022
JOURNAL OF CHEMICAL INFORMATION AND MODELING

READ 

Graph-Driven Reaction Discovery: Progress, Challenges, and Future Opportunities

Idil Ismail, Scott Habershon, *et al.*

OCTOBER 03, 2022
THE JOURNAL OF PHYSICAL CHEMISTRY A

READ 

Generative Modeling to Predict Multiple Suitable Conditions for Chemical Reactions

Youngchun Kwon, Seokho Kang, *et al.*

NOVEMBER 22, 2022
JOURNAL OF CHEMICAL INFORMATION AND MODELING

READ 

Get More Suggestions >

Supporting Information

Representing catalytic mechanisms with rule composition

Jakob L. Andersen¹, Rolf Fagerberg¹, Christoph Flamm², Walter Fontana³, Juri Kolčák^{*1,3},
Christophe V.F.P. Laurent¹, Daniel Merkle^{†1}, and Nikolai Nøjgaard¹

¹*Department of Mathematics and Computer Science, University of Southern Denmark,
Campusvej 55, 5230 Odense, Denmark*

²*Department of Theoretical Chemistry, University of Vienna,, Währinger Straße 17, 1090
Vienna, Austria*

³*Department of Systems Biology, Harvard Medical School,, 200 Longwood Avenue, 02115 Boston,
MA, USA*

A Graph Transformation for Chemistry

In this section we introduce the graph transformation concepts necessary for mathematically rigorous treatment of overlay graphs and substrate rules. While formal, the definitions utilize chemical examples. For a purely formal, category theoretical treatment, please refer to [Section B](#).

The main text offers a simplified view on the notion of a graph-transformation rule. In practice, we utilize the double pushout formalism [1, 2], in which the correspondence between atoms in L and R is established by means of a graph K and two maps from (the atoms in) K to (atoms in) L and R . Formally, a rule is a span $p = (L \xleftarrow{l} K \xrightarrow{r} R)$, where L and R are the left (in) and right (out) graphs, respectively. K is the invariant graph containing elements common to L and R as specified by the injections l and r . What is in L but not in K vanishes and what is in R but not in K appears.

The application of a rule p to a mixture G occurs in two steps, illustrated in [Figure S1A](#). First, an embedding or match, m , of L in G must be found. Second, given m , those parts of G that are in L but not in K (and thus meant to vanish) are removed to be replaced by those parts in R that are not in K (thus meant to appear), yielding the new mixture H . In the standard approach suitable for chemistry, punching a hole in G is subject to some restrictions, such as avoiding the creation of dangling bonds. Absent dangling bonds, some information is needed for filling the hole in G with R . This information comes from K and the associated maps l and r , which specify which left-over parts at the location of the match m are meant to be the same parts in both L and R . Given this identification, R can be properly “glued” in place. All this is captured by the commuting diagram shown in [Figure S1A](#), where the top row is the rule, p , and G in the bottom row is the given mixture. Punching the hole in G amounts to constructing D , which together with K provides the instruction for gluing R in place.

We write $G \xrightarrow{p,m} H$ for the transformation of G into H by rule p at the location of the match m and refer to it as a *direct derivation*. It is important to keep in mind that G and H are graphs representing fully specified molecules. Typically G comprises the reactants of a particular reaction, but G could also represent a large mixture within which a reaction takes place at the location indicated by m . (A direct

*Corresponding author: juri.kolcak@gmail.com

†Corresponding author: daniel@imada.sdu.dk




Figure S1: **A:** The example rule describing a keto-enol tautomerization, applied to a 2-hydroxyacrylaldehyde in full formalism. To hint at the actual mappings l and r , we use a spatial layout of K in which the position of an atom is meant to suggest that it is mapped to the atom at the same (relative) position in L and R . **B:** The ITS of the reaction.

derivation is really just a state transition, but we soon need more vocabulary.) If by virtue of the action of p the new graph H comes to satisfy the left pattern of rule q , we can view the application of p as enabling, or causing, an application of q . In this way, step by step, a sequence of direct derivations $G \xrightarrow{p_1, m_1} H \equiv H_1 \xrightarrow{p_2, m_2} H_2, \dots, H_{n-1} \xrightarrow{p_n, m_n} H_n$ represents a causal trajectory of transformations converting educts G into products H_n while at the same time mapping atoms in G to atoms in H_n . We refer to the whole sequence as a *derivation*.

On top of the non-bond constraints introduced in Section 2 of the main text, we need another syntactical construct specific to chemical applications. An atom that changes state, such as charge, should preserve its identity. In other words, a change of state should not be interpreted as the deletion of an atom in one state and the re-creation of an atom of the same type in the new state. We use a bullet in K to indicate that the identity of the corresponding atom is preserved but that it acquires a new label to reflect the change of state, such as going from N to N^+ . An example of such untyped, or better yet unlabeled, vertices is available in Figure S2A making use of the spontaneous breakup of a phenyl-N-(sulfonatoxy)methanimidothioate retrieved from the MetaCyc database [3].

While both non-bond constraints and unlabeled vertices in K are sound [2, 4, 5], at present, there’s a lack of comprehensive support within graph transformation frameworks. As addressing this deficiency is not our primary concern, it should be noted that our implementation circumvents the issue by exploiting the structure of our input data, and does not generalize beyond our application domain.

Note that the ITS in Figure S2B is not cyclical, even though the reaction is elementary. This is due to the changes in formal charges acting as sources and sinks for electron pairs. One can recover an ITS that is a loop by considering a virtual entity representing a free electron pair, as illustrated in Figure S2C and D. This representation would not require label changes, but relies on artificial constructs, such as the explicit transfer of the free electron pair from the sulfur atom to the oxygen atom.

A.1 Rule Composition

Our focus is on mechanisms of enzymatically catalyzed reactions. The significance of a mechanism consists in providing a series of elementary steps leading from educts to products, as well as an atom map, allowing a representation as a derivation – sequence of rule applications, each corresponding to a step of the mechanism. We thus do not aim to describe the composition of rules in full generality [6, 4], but focus on the specific case of rule composition along a derivation.

We illustrate rule composition for the base case of two steps in which first rule p_1 is applied to the initial mixture, followed by an application of rule p_2 . The case of n steps is simply an iteration of the

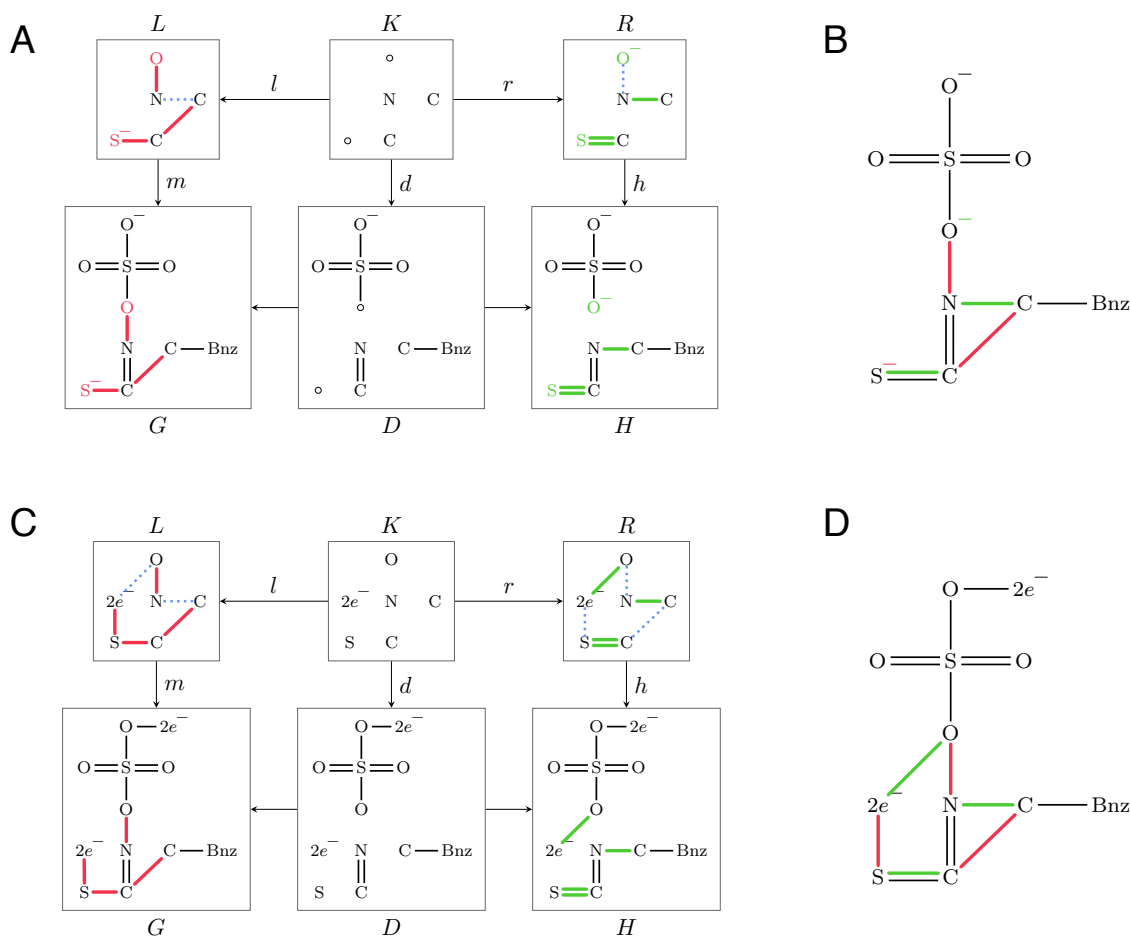


Figure S2: A rule applied to phenyl-N-(sulfonatoxy)methanimidothioate, representing its spontaneous breakup into a benzylisothiocyanate and a sulfate. This rule modifies formal charges, requiring the affected atoms to appear untyped in the invariant graph K . The ITS of the reaction is depicted on the right hand side.

base case. Conceptually, the left pattern L of the composite rule p must collect the conditions (the graph parts) demanded by the left pattern of p_1 plus those conditions of the left pattern of p_2 that have not been created by the action of p_1 . Likewise, the right pattern R of p must contain the right pattern of p_2 plus those parts of the right pattern of p_1 that have not been modified by p_2 . The conversion of this specification into a formal construction makes use of standard operations on graphs, which we briefly explain to provide a self-contained presentation using as an example the first two steps of the mechanism in the main text, Figure 2.

The key operation is the gluing together of two graphs, which occurs by declaring certain nodes and links in one graph as equivalent to certain nodes and links in the other, then identifying (merging) the equivalent graph elements to combine the two graphs. The declaration of equivalence amounts to a “gluing instruction”. It must respect type; so we cannot identify an O with a C or a single bond with a double bond.

For example, the graph I at the top of Figure S3 is a gluing instruction asserting that the $O^- - C$ part is the same in both R_1 and L_2 ; likewise for $H - N^+$. This correspondence is made explicit by the mappings i_1 and i_2 , which are given here implicitly by the layout of the atoms. Note that in this example R_1 and L_2 are isomorphic, which illustrates why a gluing instruction is necessary: the fact that I does




Figure S3: Diagram showing the construction of the graph H as detailed in the text. Note that R_1 and L_2 look the same in this case. Yet, the gluing instruction I does not mention the neutral oxygen atoms, which means that their instances in R_1 and L_2 are not equivalent and thus not be identified in the merging of R_1 and L_2 . H can be viewed as a union of R_1 and L_2 based on the identifications made by I with the maps i_1 and i_2 . In this example R_1 and L_2 are the right and left patterns of the rules representing the first and second steps, respectively, of the triacylglycerol lipase mechanism. The source of information captured in I and the maps i_1 and i_2 is the gray graph H_1 in the middle, which is the fully specified molecular graph (the mixture) just after the first step has been applied (with h_1 showing the location of R_1) and prior to the second step, showing the location m_2 where it will be applied. This information is part of “knowing the mechanism” and distinguishes the two C–O components. The mechanism is necessary to define the gluing instruction I .

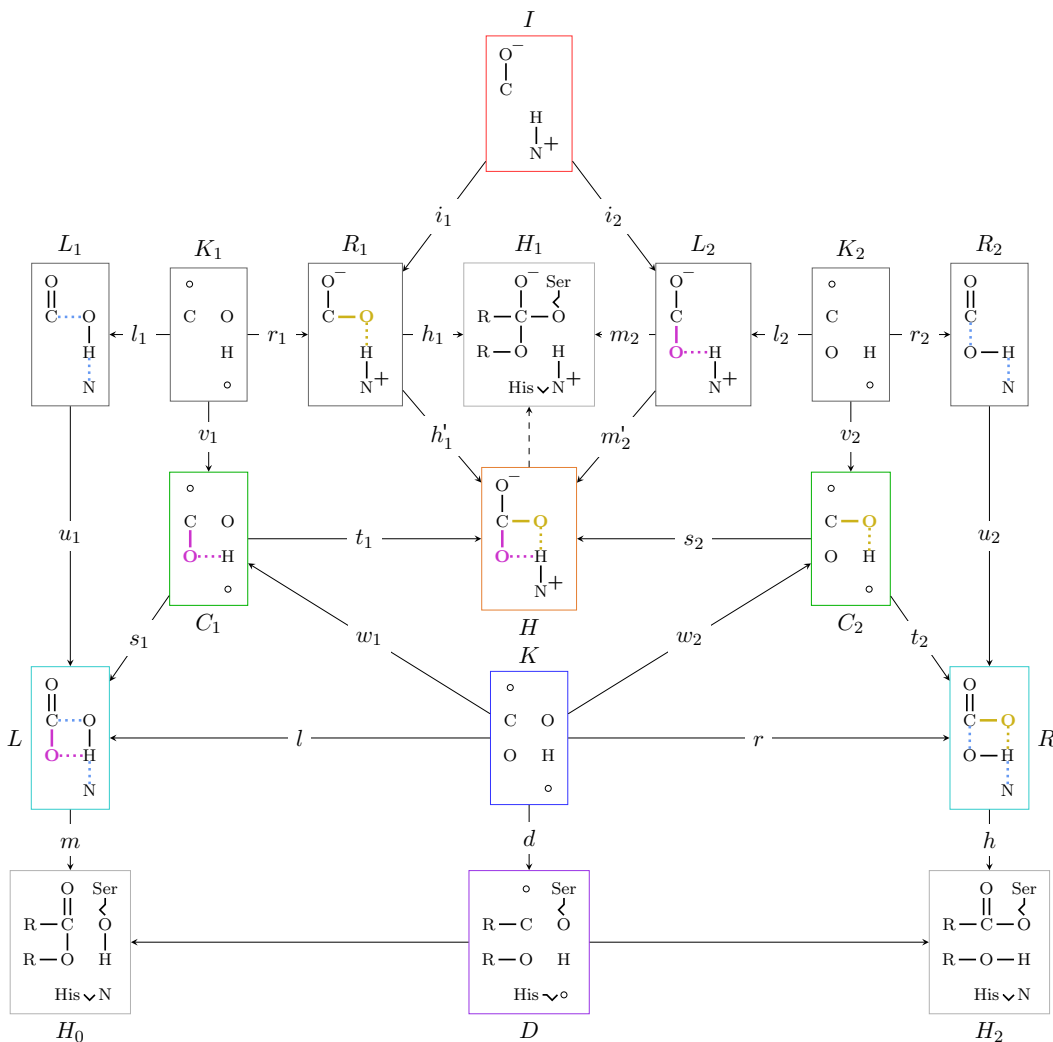


Figure S4: Diagram with the sequence of constructions needed to compose the two rules $p_1 = (L_1 \xleftarrow{l_1} K_1 \xrightarrow{r_1} R_1)$ and $p_2 = (L_2 \xleftarrow{l_2} K_2 \xrightarrow{r_2} R_2)$. The information in I comes from the mechanism, i.e. the result (h_1) of the application of p_1 in a concrete mixture and the embedding (m_2) of the left pattern of p_2 in preparation for the subsequent application of p_2 (gray structures). This information allows the construction of H , which contains the graphical elements that must be merged with L_1 and R_2 to obtain the left and right patterns L and R , respectively, of the composite rule. After H , C_1 (C_2) is constructed, which in turn allows the construction of L (R) as detailed in the main text. The commutative property underlying these manipulations enables their automatic execution.

not mention a neutral O means that the (uncharged) oxygen atoms in R_1 and L_2 are not to be identified; they are distinct instances of O. Given I and the maps i_1 and i_2 , the graphs R_1 and L_2 can be glued together uniquely yielding H . If R_1 and L_2 contained multiple instances of, say, $O^- - C$, the mappings i_1 and i_2 would determine which one corresponds to which. This raises the question about the choice of i_1 and i_2 when there are multiple possibilities. The answer comes from knowing the actual mechanism, given by the execution of the rules in the concrete setting of a mixture of reactants. The notion of mechanism includes the graph H_1 , which is the intermediate mixture, along with the embedding maps h_1 and m_2 telling us where rule p_1 acted and where rule p_2 is supposed to act. In the present case, the mechanism distinguishes the two (uncharged) oxygen atoms. The O mentioned in R_1 comes from the serine, which binds the substrate to the enzyme as a result of step 1; whereas the O in the isomorphic L_2 belongs to the substrate, from which it will be severed in step 2. Hence, the composition of two rules, $p_1 = (L_1 \xleftarrow{l_1} K_1 \xrightarrow{r_1} R_1)$ and $p_2 = (L_2 \xleftarrow{l_2} K_2 \xrightarrow{r_2} R_2)$ depends on where in the mixture p_1 and p_2 act. This location is part of the mechanism and is given by H_1 , which contains the result of p_1 (i.e. R_1) and satisfies the requirements of p_2 (i.e. L_2). We thus annotate the composition operator with the location information: $p = p_1 \bullet_{H_1} p_2$.

We next use the gluing construction to determine the left (L) and right (R) pattern of the composite rule, as illustrated in [Figure S4](#). The orange-framed graph is H of [Figure S3](#), which is the gluing of R_1 —the structure created by p_1 —with L_2 —the structure required by p_2 . Thus, H contains both the graph fragments requested by L_2 but not delivered by p_1 and the modifications caused by p_1 that survive the action of p_2 (because their presence is ignored by L_2). We only need to tease these contributions apart and glue them to L_1 and R_2 to obtain L and R of the composite rule.

Note that H can also be viewed as the gluing of R_1 with an unknown graph X using the gluing instructions provided by K_1 . The difference between R_1 and H , highlighted in magenta, consists of the fragment $[-O\cdots]$ comprising the oxygen atom of the substrate, its single bond to the carbon of the substrate and the non-bond constraint with the H of the protonated histidine. This fragment must therefore occur in X . The unknown graph we seek is given by C_1 , which is H from which R_1 has been removed while keeping the fragments indicated by K_1 . With C_1 in hand, we can determine L as the gluing of L_1 with C_1 . We now see that the $[-O\cdots]$ fragment, which is part of the structural condition required by p_2 and not delivered by p_1 , has been included as a requirement in the left pattern L of the composite rule. Using the same reasoning we transport the other $[-O\cdots]$ fragment, highlighted in yellow, to the R -side. This fragment comprises the oxygen of the serine, its single bond to the carbon of the substrate and the non-bond constraint with the protonated histidine. The same logic is deployed to obtain the invariant graph K of the composite rule as the gluing instruction that directs C_1 and C_2 to be combined into a graph compatible with H .

The composite rule $p = (L \xleftarrow{l} K \xrightarrow{r} R)$ can be applied to a mixture H_0 in which it induces the direct derivation (the state change) $H_0 \xrightarrow{p,m} H_2$, as shown at the bottom of [Figure S4](#), without a need for the intermediate stage H_1 . This procedure can now be iterated for all steps of a mechanism to construct the composite rule of the overall reaction. For the mechanism from main text [Figure 2](#), this composite rule p can be found as $p = (L \xleftarrow{l} K \xrightarrow{r} R)$ in the second row of [Figure S5](#).

We can now see that the partial catalog of changes the rule composition retains in place of the full causal information in the mechanism is largely recorded withing the invariant graph K of the composite rule. (i) The doubly-bound O and the N must have undergone a transient change, because K has no type specification at those nodes and yet they are present in both L and R . (ii) Likewise, the double bond is in both L and R , but not in K ; hence, it must have been transiently destroyed and then reconstituted. (iii) The N and H must have been temporarily bound to each other, because the non-bond constraint is present in L and R , but not in K . (iv) One C–O bond has been cleaved and a different one has been created. (v) Similarly, one O–H bond was lost and another was gained. The latter two changes are made explicit as differences between L and R . The point is that the composite rule also preserves transitory

changes, because, by construction, K no longer specifies just the parts that are the same in L and R , but rather the parts that stay the same *throughout* the mechanism.

The composite rule represents the reaction mechanism in a compact form. The combined graph, or OG then contains all atoms, bonds and explicit non-bond constraints that appear on either side of the composite rule. Since the rule also records transient changes, we use a color scheme to annotate the change of type that various parts experience as they pass from L to R . A bond (or atom) is shown in green when the net change across the reaction is a formation and in red when the net change is an elimination. Bonds and atoms whose original state is eliminated and then re-introduced are shown in purple. For non-bond constraints, which flag an absence, we color transient changes in blue to emphasize that an introduction is followed by an elimination (the opposite of purple transients).

The example used to illustrate the process of rule composition included steps that covered only the reaction center, that is, those parts of a molecular graph that are subject to modifications. This is why the invariant graph K does not show any parts that remain constant throughout the mechanism, except for atoms. In practice, a rule may include molecular graph fragments that are known to constitute a condition for reaction but are not altered by the mechanism, such as heteroatoms or functional groups that modify the electronic density or substituents that act as steric constraints. These truly constant parts (as opposed to the transiently-modified “constant” parts) show up in L , R and in K . They are inherited by K because they would occur in K_1 and K_2 , and thus in C_1 and C_2 from which K is constructed.

A.2 Substrate Rules

Finally, we present the formal version of the substrate rule inference. To construct the substrate rule associated with a catalytic process, we proceed as illustrated in [Figure S5](#). In essence, we need to remove the parts representing the catalyst from the composite rule $L \xleftarrow{l} K \xrightarrow{r} R$, shown in the second row of [Figure S5](#). (Note that the composite rule is not the one of [Figure S4](#), which resulted from composing only step 1 and step 2, but the rule resulting from the composition of all four steps of the full mechanism of the main text [Figure 2](#).) The graph C is our choice of helper molecules. The graph C , alongside the mappings c_1 and c_2 , identifies the catalyst in the educt mixture E and the product mixture P . By using the empty graph as a gluing instruction we effectively remove the graph parts of C from E and P to yield E' and P' (top row), respectively. L' is the intersection of L and E' ; similarly for R' . To obtain the graph K' for the substrate rule we first construct N_1 , which is the intersection of K (whose map l tells us what stays invariant in L) and L' , since L' is in L . Similarly for the R' side to get N_2 . K' then is the intersection of N_1 with N_2 .

The choice of C , c_1 and c_2 , is subject to to the obvious constraint that C must be a subgraph of both E and P . The construction of the substrate rule and the applicability of the rule, however, pose further constraints. In particular, the constraints arise due to the possibility of introducing dangling bonds. The potential to introduce dangling bonds is twofold. During the construction of E' and P' , and during the application of the substrate rule to the original mixture.

The case of E' and P' is relatively simple. Any vertex in C is being removed from E to obtain E' . All edges such a vertex has in E have to be reflected in C , including their endpoints. The situation is the same for P' and P , constraining C to correspond to connected components, or fully specified molecules, in E and P . This constraint is in general not an issue for the substrate rules, as they are constructed by explicitly removing fully specified molecules. However, as we demonstrated with the notion of catalytic complement OG, one may be interested in considering only specific functional groups (“catalytic sites”) for inclusion in the graph C .

For example, recall the lower left quadrant in [Figure 8A](#) of the main text, which depicts an overlay graph of M-CSA [entry 748](#). In this case, the electron source is a nitrogen of an amine group that belongs to the substrate. The intra-molecular connection between the carboxylic ester carbon and the catalytic amine group consists of a single carbon atom, [Figure S6](#), but is abstracted away (as symbolized by the




Figure S5: Diagram for constructing the substrate rule. The colors indicate the progression of graph constructions: red first, then orange, then green and finally purple. See text for details.

way line in the main text Figure 8A) because our minimalistic rules and OGs only include bonds and atoms that directly participate in the reaction. The catalytic nitrogen of the substrate thus appears as a disconnected graphical component. In case the connecting carbon is included in the rule and the OG, one has to ensure that the C–N bond to the amine nitrogen is removed during the construction of the catalytic complement OG, alongside the nitrogen itself. The C–N bond would then necessarily appear in the graph C . Preventing the erasure of the carbon atom from E' and P' is achieved by adding it to the empty graph at the top of Figure S5.

The second instance occurs after the substrate rule has been obtained and restricts its applicability to the original mixture. The restriction becomes relevant if graph L' has more vertices than K' , meaning that the substrate rule deletes atoms, which have to show up outside the substrate in the catalytic part of the reaction. These “catalytic atoms” are identified by choosing the maps c_1 and c_2 , which define an atom map from the catalytic parts of E to the catalytic parts of P . However, the atom map defined by c_1 and c_2 could be at odds with the atom map defined by the mechanism and represented by the morphisms l and r of the associated composite rule. This could thwart the applicability of the substrate rule. To avoid this situation, the following must be ensured: all bonds of an atom in E that is mapped by c_1 and c_2 to an atom in P different than the atom to which it is mapped by l and r must be represented in L . In other words, if an atom that is part of a catalytic moiety becomes part of another catalytic moiety (hence is being created or deleted by the substrate rule), then all of its bonds have to be accounted for in the composite rule. This can usually be achieved by choosing c_1 and c_2 in a way that maximizes their overlap




Figure S6: The amine group of the tRNA bound D-amino acid substrate, acting as a catalyst of (M-CSA entry 748), in relation to the carboxylic ester carbon (black).

with the atom map of the mechanism. There are interesting exceptions, however, which we explore in greater detail in Section E.

B Graph Transformation in Category Theory

This material reiterates the constructions of the composite rule and substrate rule from the previous section in full categorical formalism. At least an intuitive understanding of the categorical notions of span, pullback, pushout and pushout complement is needed.

First, we revisit the composition of two rules $p_i = L_i \xleftarrow{l_i} K_i \xrightarrow{r_i} R_i$ equipped with direct derivations $H_{i-1} \xrightarrow{p_i, m_i} H_i$ for $i \in \{1, 2\}$. Figure S7 shows the construction of the shared subgraph H via a pullback from the co-span $R_1 \xrightarrow{h_1} H_1 \xleftarrow{m_2} L_2$ and a subsequent pushout. Note that h_1 and m_2 come from the direct derivations as illustrated in Figure S1A.




Figure S7: Commutative diagram showing the construction of the graph H , as a pullback from the co-span $R_1 \xrightarrow{h_1} H_1 \xleftarrow{m_2} L_2$ and a subsequent pushout of the obtained span $R_1 \xleftarrow{i_1} I \xrightarrow{i_2} L_2$. The dashed arrow indicates the unique commuting morphism embedding H into H_1 .

The graph H (which is part of “knowing the mechanism”) is an input to the rule composition. The composition itself is carried out in three steps, Figure S8:

1. C_1 and C_2 are constructed as pushout complements of $K_1 \xrightarrow{r_1} R_1 \xrightarrow{h'_1} H$ and $K_2 \xrightarrow{l_2} L_2 \xrightarrow{m'_2} H$, respectively.
2. L and R are obtained as pushouts of the spans $L_1 \xleftarrow{l_1} K_1 \xrightarrow{v_1} C_1$ and $R_2 \xleftarrow{r_2} K_2 \xrightarrow{v_2} C_2$.

3. The graph K is produced as a pullback of the co-span $C_1 \xrightarrow{t_1} H \xleftarrow{s_2} C_2$.

The rule $p = L \xleftarrow{l} K \xrightarrow{r} R$ is completed by taking $l = s_1 \circ w_1$ and $r = t_2 \circ w_2$. Finally, the composite derivation $H_0 \xrightarrow{p,m} H_2$ is such that $m \circ u_1$ commutes with m_1 and $h \circ u_2$ commutes with h_2 .




Figure S8: Commutative diagram for rule composition of two rules $p_1 = (L_1 \xleftarrow{l_1} K_1 \xrightarrow{r_1} R_1)$ and $p_2 = (L_2 \xleftarrow{l_2} K_2 \xrightarrow{r_2} R_2)$ with direct derivations $H_0 \xrightarrow{p_1, m_1} H_1$ and $H_1 \xrightarrow{p_2, m_2} H_2$ respectively. The resulting rule is $p = p_1 \bullet_{H_1} p_2 = (L \xleftarrow{l} K \xrightarrow{r} R)$ with $l = s_1 \circ w_1$ and $r = t_2 \circ w_2$. The numbers in parentheses denote the order in which the graphs are constructed by completing the commuting (sub)diagram.

The rule composition yielding also a direct derivation allows us to chain the compositions until the entire mechanism is represented in a single, composite rule.

To obtain the substrate rule, we once again rely on the categorical constructions of pullback and pushout complement. Given the monomorphisms $C \xrightarrow{c_1} E$ and $C \xrightarrow{c_2} P$, the construction of the substrate rule (Figure S9) is the following multi-step process.

1. The complement of C in E and the complement of C in P , E' and P' respectively, are formally obtained as pushout complements of $\emptyset \rightarrow C \xrightarrow{c_1} E$, respectively $\emptyset \rightarrow C \xrightarrow{c_2} P$, where \emptyset is the empty graph.
2. The graphs L' and R' are constructed as pullbacks of the co-spans $E' \rightarrow E \xleftarrow{m} L$ and $P' \rightarrow P \xleftarrow{h} R$, effectively projecting L and R on the substrate subgraphs of E and P , respectively.
3. In the third step, two helper structures, N_1 and N_2 , are obtained as pullbacks of the co-spans $L' \rightarrow L \xleftarrow{l} K$ and $R' \rightarrow R \xleftarrow{r} K$, respectively.
4. Finally, K' is constructed as a pullback of the co-span $N_1 \rightarrow K \leftarrow N_2$.

The substrate rule $L' \xleftarrow{l'} K' \xrightarrow{r'} R'$ is finalized by taking $l' = u_1 \circ n_1$ and $r' = u_2 \circ n_2$. The diagram in Figure S9 depicts the construction of the substrate rule for a given C , c_1 and c_2 .




Figure S9: Commutative diagram inferring the substrate rule using categorical constructions. E' and P' are the substrate part of the educt graph and product graph, respectively, constructed as pushout complements. The left and right graphs L' and R' of the substrate rule correspond to the shared part of E' and L and of P' and R , respectively, constructed as pullbacks. Similarly, K' corresponds to the shared part of L' , R' and K . The numbers in parentheses denote the order in which the graphs are constructed.

C Atom Exchange between Substrate and Catalyst




Figure S10: An overlay graph of a formyl-CoA transferase reaction (M-CSA [entry 155](#)). The aspartate of the catalyst gives away its O^- to the formate product, indicated by the red bond at the bottom left. The aspartate is subsequently restored by an oxygen coming from the oxalate educt, green bond at the bottom right.

In the main text we mention the possibility that substrates and catalytic sites trade atoms. While this is very common for hydrogen atoms, as in the running example of the main text (M-CSA [entry 218](#)),

heavier atoms, like oxygen, can also be exchanged, as shown in [Figure S10](#).

D Limits of the Rule Composition Approach

Rule composition amounts to coarse-graining a mechanism, because it cannot retain all the information contained in the ordered sequence of steps that constitutes a mechanism. Here we summarize what kind of information is lost by using as illustrative case the D-alanine transaminase reaction (M-CSA [entry 66](#)) whose automatically generated overlay graph is depicted in [Figure S11A](#). (The layout, however, was done by hand.)

The reaction mechanism extends over 12 steps and uses pyridoxal phosphate (PLP) as a cofactor. PLP starts out bound as a Schiff base to the amino end of a lysyl residue of the enzyme. It functions as both an electron sink through the positively charged nitrogen atom in its ring and as a transient carrier of an amino group, which it accepts from the glutamate substrate and transfers to the pyruvate with the assistance of an H_2O molecule. In this process, among other events, a bond between an H and the N of the substrate amino group is broken and later reconstituted (the purple N–H edge in the overlay graph). In another step an electron pair is transiently parked at the ring nitrogen of PLP, which entails transiently flipping a number of double and single bonds (the purple chain $\text{C}-\text{C}=\text{C}-\text{C}=\text{N}^+$ in the overlay graph in [Figure S11A](#)).

We observe that the overlay graph of the composite rule correctly records the N–H edge as transiently broken, but it also records the C–C bonds in the purple chain as transiently broken, when, in fact, they transiently change from single to double bonds; similarly for the C=C bonds in that same chain. Moreover, while it is the case that the $\text{C}=\text{N}^+$ bond between PLP and lysine is transiently broken, as reported in the overlay graph, it is also the case that in the process the same C and N are transiently linked through a single bond. The overlay graph only reports the “outer” or top-level modification and is blind to transient modifications nested within. In the same vein, the overlay graph does not record multiple transient modifications of the same kind between two atoms.

A non-bond constraint obviously carries no bond type information. The overlay graph can, therefore, only record a transient bond formation between PLP and the N of the substrate (the blue C–N edge), but cannot report that at some point it also is a double and not only a single bond.

Finally, the overlay graph does not record transient changes that precede permanent changes, whether creation or dissolution, of a bond. For example, the C–N single bond of the glutamate first becomes transiently a double bond before being broken; similarly, the newly created C–N bond of the alanine product, resulting from the amino transfer to pyruvate, is first created as a double bond before being reduced to a single bond.

All this information is lost because nodes and edges of the invariant graph K of a rule can only be typed by a label that matches these nodes and edges in both L and R . Thus, if an edge is transiently modified from single bond to double bond, it shows up as a single bond in the L and R pattern of the composite rule, but as absent in the invariant graph K . From this, one can only deduce the existence of a transient modification, but not its kind. In the same vein, one cannot use K to infer the multiplicity of the transient modifications, shall the same one occur multiple times, nor quantify or even detect that such an information loss occurred during the composition.

For comparison, [Figure S11B](#) shows the overlay graph annotated by hand to aid in locating the incomplete information reported in panel A. It is possible to add type distinctions of the kind we used for state changes of atoms, but at that point one might as well forgo coarse-graining by dropping the composition construction altogether and instead work with the detailed sequence of mechanistic steps. This, on the other hand, might provide too much detail, limiting insight and hampering the automation of search. Especially with regard to the latter, the information lost by rule composition in no way limits retrieving reactions that are compatible with the mechanism implicit in the composite rule, which still accurately records the graphical requirements for a reaction, given its mechanism.




Figure S11: **A:** Overlay graph of the D-alanine transaminase reaction (M-CSA [entry 66](#)) resulting from rule composition. **B:** An overlay graph of the same reaction created by hand to illustrate the full extent of information contained in the detailed mechanism.

E Substrate Rules and the Notion of Catalyst

The substrate rule is derived from the composite rule by removing the parts that are responsible for catalysis in the reaction. The idea behind the substrate rule is to retain the condition of applicability of the mechanism while being as agnostic as possible about the chemical realization of the catalyst. In ultimate analysis, however, the idea of the substrate rule must be viewed as a heuristic, as it can infringe on the limits of the graph transformation framework. We briefly clarify this issue with two examples.

The first example is an isochorismate synthase reaction (M-CSA [entry 325](#)) in which the respective OH of a water molecule becomes a new alcohol group of the substrate and an original alcohol group of the substrate is released as a water molecule ([Figure S12](#)). According to the mechanism, the incoming and the outgoing water molecule share one hydrogen atom. This hydrogen atom is abstracted from the incoming water molecule, parked at a lysine residue and retrieved later in forming the outgoing water molecule during a clean-up step that restores the enzymatic state. The other hydrogen atoms of both the incoming and outgoing H_2O are the ones that travel with the OH components. These hydrogen atoms never actively participate in the reaction (gray) and are therefore not represented in the left pattern of the composite rule or the associated substrate rule. When applying the substrate rule to a mixture, we need to be able to delete the oxygen atom indicated by an arrow in [Figure S12A](#) from the substrate graph. Yet, there is no way of specifying this deletion within the double pushout framework, unless the bond to the passive hydrogen is included in the composite (and substrate) rule. Thus, the substrate rule cannot be applied despite there being only a single water molecule on the left and the right side of the reaction, qualifying it as a catalyst and making the choice of c_1 and c_2 in [Diagram S9](#) unique. The key is that the H_2O on the left and the right consist of different atom identities, which makes the catalytic atom map defined by c_1 and c_2 (in which the H_2O on the left and the right appear identical) jar with the atom map of the mechanism.




Figure S12: **A:** An overlay graph of an isochorismate synthase reaction (M-CSA [entry 325](#)). **B:** The substrate overlay graph of the same reaction with H_2O in C . The rule is not applicable to the original educt mixture as the hydrogen bond (gray) of the red oxygen atom prevents its deletion.

As a result, and for our purposes, the water molecule(s) cannot be declared catalytic in the sense of belonging to graph *C* in [Diagram S9](#), unless the inactive hydrogen atoms are explicitly added to the invariant graph of the composite rule, thereby enabling their deletion and creation alongside the respective oxygen atoms.

The change of identity of catalytic atoms during a reaction need not by itself lead to an inapplicable substrate rule. Consider, for example, the overlay graph in [Figure S11A](#), where the oxygen atom of a water molecule becomes an acyl group, and an acyl group of an educt is released as a water molecule. Since both hydrogen atoms of the water molecules are active participants in this reaction, they become part of the composite rule and there is no difficulty with including H₂O into the catalytic graph *C* when constructing the substrate rule. Hence the issue is for the composite rule to include “enough” context. If the context is minimal, because the rules only refer to the reaction center as they do in the present paper, inapplicable substrate rules can arise. The modeler is, of course, at liberty to add more context to the composite rule so as to make the substrate rule applicable, but this then might be dictated more by mathematical than chemical necessity.




Figure S13: An overlay graph of a 3-dehydroquinate dehydratase reaction (M-CSA [entry 55](#)). Two water molecules, framed by brown boxes, occur in the reaction. The water molecule labeled 1 is catalytic, serving as a proton relay between the asparagine and the substrate. The other, labeled 2, is a product formed by elimination of the alcohol group from the substrate.

Another interesting case is provided by M-CSA [entry 55](#), which describes a reaction that dehydrates 3-dehydroquinate to 3-dehydroshikimate. An H₂O molecule must be supplied as an educt because it is needed as a catalyst in the first two steps of the reaction (where it twice establishes a hydrogen relay). A new H₂O molecule is created in this process by removing an OH group from the substrate and combining it with a hydrogen from a histidine residue of the enzyme. At this point the mechanism has yielded a H₂O and 3-dehydroshikimate, in addition to the original H₂O molecule. However, the chemical state of the enzyme still needs to be restored. This is accomplished in the final step by using a H₂O molecule

once more as a hydrogen relay. In terms of the mixture contents in that phase of the mechanism, the second catalytic deployment of water might involve the just newly generated H₂O molecule, rather than the one present initially. Overall, the reaction starts out with an H₂O and 3-dehydroquinone to yield 2H₂O and 3-dehydroshikimate. The overlay graph of the composite rule is shown in [Figure S13](#), in which we identify and label the water molecules.

When constructing the substrate rule, the question arises as to who are the catalysts in this reaction. Clearly, the enzyme is a catalyst, but so is the H₂O #1 which catalyses the first two steps of the mechanism and can obviously also catalyze the third clean-up step; but so could be a combination of H₂O #1 and H₂O #2. When choosing the amino acid residues and H₂O #1 as the catalytic graph C , we obtain an applicable substrate rule. In this case, both c_1 and c_2 identify the same H₂O #1 on the left and the right of the rule, thus automatically maximizing the compatibility with the atom map of the mechanism. If c_2 were to point at H₂O #2 instead, we would run into the same applicability issue discussed in the context of [Figure S12B](#).

That said, H₂O #2 clearly can be a catalyst in the third step of the mechanism. From a network perspective, H₂O #2 is therefore both a product and a catalyst. Indeed, it makes the reaction autocatalytic (albeit in a likely uninteresting way, given the abundance of water). Yet, the substrate rule approach is unable to capture H₂O #2 as a catalyst, since it does not appear on the educt side of the overall reaction.

References

- [1] A. Habel, J. Müller, and D. Plump, “Double-pushout graph transformation revisited,” *Mathematical Structures in Computer Science*, vol. 11, no. 5, pp. 637–688, 2001.
- [2] J. L. Andersen, C. Flamm, D. Merkle, and P. F. Stadler, “A Software Package for Chemically Inspired Graph Transformation,” in *Graph Transformation. ICGT 2016. Lecture Notes in Computer Science* (R. Echahed and M. Minas, eds.), vol. 9761, pp. 73–88, Cham, Switzerland: Springer, 2016.
- [3] R. Caspi, R. Billington, I. M. Keseler, A. Kothari, M. Krummenacker, P. E. Midford, W. K. Ong, S. Paley, P. Subhraveti, and P. D. Karp, “The MetaCyc Database of Metabolic Pathways and Enzymes - a 2019 Update,” *Nucleic Acids Research*, vol. 48, pp. D445–D453, 10 2019.
- [4] J. L. Andersen, C. Flamm, D. Merkle, and P. F. Stadler, “Rule composition in graph transformation models of chemical reactions,” *MATCH, Communications in Mathematical and in Computer Chemistry*, vol. 80, no. 3, pp. 661–704, 2018.
- [5] N. Behr, J. Krivine, J. L. Andersen, and D. Merkle, “Rewriting theory for the life sciences: A unifying theory of ctmc semantics,” *Theoretical Computer Science*, vol. 884, pp. 68–115, 2021. TR: <https://arxiv.org/abs/2106.02573>.
- [6] J. L. Andersen, C. Flamm, D. Merkle, and P. F. Stadler, “Inferring Chemical Reaction Patterns Using Rule Composition in Graph Grammars,” *J. Syst. Chem.*, vol. 4, no. 1, p. 4, 2013.