Published in final edited form as:

J Chem Inf Model. 2009 September ; 49(9): 2034–2043. doi:10.1021/ci900157k.

No electron left behind: a rule-based expert system to predict chemical reactions and reaction mechanisms

Jonathan H. Chen and Pierre Baldi^{*,#}

Institute for Genomics and Bioinformatics and Department of Computer Science, School of Information and Computer Sciences, University of California, Irvine, Irvine, CA 92697-3435

Abstract

Predicting the course and major products of arbitrary reactions is a fundamental problem in chemistry, one that chemists must address in a variety of tasks ranging from synthesis design to reaction discovery. Described here is an expert system to predict organic chemical reactions based on a knowledge base of over 1,500 manually composed reaction transformation rules. Novel rule extensions are introduced to enable robust predictions and describe detailed reaction mechanisms at the level of electron flows in elementary reaction steps, ensuring that all reactions are properly balanced and atom-mapped. The core reaction prediction functionalities of this expert system are illustrated with applications including: (1) prediction of detailed reaction mechanisms; (2) computer-based learning in organic chemistry; (3) retro synthetic analysis; and (4) combinatorial library design. Select applications available via <http://cdb.ics.uci.edu>.

1 Introduction

Among the most fundamental problems in organic chemistry is predicting the course and major products of arbitrary reactions. In addition to being a fundamental scientific problem, reaction prediction is also important for several practical applications including the planning of new chemical experiments and syntheses. Seminal work in computer-aided reaction prediction was achieved with the CAMEO¹ and EROS² systems and several other projects have made their own advances (e.g., Beppe³, ROBIA⁴, SOPHIA⁵, ToyChem⁶), however most computer reaction prediction systems have fallen out of support over time. Thus developing an expert system capable of reliable reaction predictions remains one of the most important and unsolved problems in chemoinformatics^{7, 8}.

The relative lack of emphasis and support for reaction prediction is surprising given its fundamental importance for organic chemistry, especially considering the amount of attention given to the complementary problem of retro synthesis. Although these two problems are closely intertwined, historically more attention has been given to computer-aided retro synthetic analysis⁹, where one wishes to identify a synthetic pathway to yield a desired target product. A likely reason for this imbalance is the more obvious relevance of retro synthesis towards obtaining important small molecules, including the majority of pharmaceutical drugs and natural products. Even within the scope of retro synthetic analysis however, reaction prediction is of direct relevance to solving one of the two key components of the analysis problem. The first component of the problem is the generation of retro synthetic suggestions while the second component is the validation of these suggestions as viable synthetic reactions. Without consideration for reactivity issues in the second component, generating retro synthetic

*Corresponding author. pfbaldi@ics.uci.edu.

#Department of Biological Chemistry, University of California, Irvine.

suggestions is relatively straightforward. A common approach involves searching a database of reactions or transformation rules for reaction centers that match the target compound of interest and proposing analogous transformations. Figure 1 illustrates how such suggestions, based on analogous examples, often do not consider functional group compatibilities and other unexpected reactivity issues that will invalidate the proposed reaction.

Existing computer-aided synthesis design systems have each addressed this problem of interfering chemical functionality to different degrees. The classic solution is to add “exclusion rules” to the suggested transformations. For the example in Figure 1, an exclusion rule could be added stating that this organometallic addition should only be suggested if none of the participating molecules contains an OH group. However, the problem is more complex because there are many other exclusion rules that would also be necessary in this example, such as the absence of SH, NH, other carbonyl, or nitrile groups. A more versatile option that has the potential to completely solve this problem is to develop a robust reaction predictor that can foresee these unexpected side reactions. To address the reaction validation component of retro synthetic analysis, a reaction predictor could simply execute a virtual reaction on any proposed precursors to verify that the intended target is actually produced.

Beyond the scope of supporting retro synthetic analysis, a robust reaction predictor would have many other immediate applications. For example, a reaction predictor could: (1) systematically generate many reactions to power combinatorial library design and development¹⁰; (2) dynamically generate and validate content to support chemical education¹¹; (3) propose mechanisms to explain the course of a reaction^{12, 13}; and (4) reveal previously undiscovered and useful reactivity.

2 Methods

2.1 System Overview

We have developed a reaction expert system to predict the major products of a reaction, given a combination of starting materials and reagents. This functionality is implemented through two primary modules, a knowledge base of transformation rules and an inference engine to process those rules (Figure 3).

A key design decision for the system is determining what the knowledge base of transformation rules represents, and in particular, at what level of detail does the system model the predicted reactions. Most past systems have used a knowledge base of transformation rules that reflect the overall reactions from starting materials to final products (Figure 2a). However, using a single rule to reflect an overall “macroscopic” reaction obscures the “microscopic” elementary steps that underlie multi-step reaction mechanisms (Figure 2b). To capture this mechanistic detail, the individual rules in our system are instead designed to mirror elementary reaction steps, from which the “macroscopic” reactions can be derived.

While the system’s transformation rules model reactions at the level of elementary processes, users are typically not interested in directly observing this level of detail. Instead, users typically prefer interacting at the level of overall reactions or even more broadly at the level of general reagents and reaction conditions. To accommodate this high level interaction, the detailed transformation rules are aggregated into reagent models that represent general chemical reagents and reaction conditions (e.g., hydrobromic acid), which can then predict the overall course of specific reactions (e.g., alkene hydrobromination). Furthermore, to develop richer and more robust predictions, the elementary transformation rules are extended with additional information and control logic such as mechanistic electron flow specifications and priority values.

2.2 Elementary Transformation Rules

The core elementary rules in the system describe chemical structure transformations using the SMIRKS language, a simple extension of the SMILES (molecule) and SMARTS (chemical pattern matching) languages¹⁴, which is processed using the OEChem toolkit¹⁵ from OpenEye Scientific Software. Though the SMIRKS specification does not require it, all the reaction equations represented by the transformation rules in the system are fully balanced with reactant atoms precisely mapped to corresponding product atoms. Ensuring that all reaction equations are fully balanced and atom-mapped is a detail often neglected by chemical data systems and even human chemists, but it is critical to ensure that transformation rules model elementary reaction steps rigorously. Table 1 lists examples of SMIRKS transformation rules that correspond to the elementary steps of the reaction mechanism depicted in Figure 2b. Currently over 1,500 distinct transformation rules have been manually composed in our system.

2.2.1 Electron Flow Specifications—The reaction transformation rules developed for this expert system are designed to mirror elementary reaction steps, which makes it relatively straightforward to extend their function to generating curved arrow mechanism diagrams^{12, 13}. This is achieved by attaching to each elementary transformation an additional string indicating where the flow of electrons should begin and end within the reaction intermediates. Figures 4a and 4b illustrate this method by applying a SMIRKS transformation rule to predict the product of an elementary step in combination with an electron flow specification.

The electron flow specification language, described below and illustrated in Figures 4a and 4b, was created for this reaction expert system as a SMIRKS language extension to support mechanistic detail in reaction transformation rules. The typical form of one of these specifications is “ $n_1, n_2=n_3, n_4$ ” where n_1, n_2 are the indexes associated with the source atoms flanking the bond of origin for the electron flow arrow while n_3, n_4 are the indexes associated with the target atoms flanking the new bond that will be formed by the elementary reaction step. A similar string like “ n_1, n_2-n_3, n_4 ” represents the movement of a single electron (i.e., a free radical reaction) instead of the more typical movement of a pair of electrons. The complete set of symbols used in this language is listed in Table 2.

While using this specification language, certain nuances in electron arrow pushing diagrams must be highlighted. One potential issue is that the specification may seem to imply that arrows can originate from the nuclei of atoms when in reality they are meant to represent the movement of the electrons. Obviously, the intended meaning in these scenarios is that the arrows represent the movement of the electrons (lone pair or free radical) associated with the atom, and not of the actual atom nucleus. Thus the specification language assumes that the user is capable of identifying lone pair and free radical electrons. Unfortunately, ChemAxon's MarvinView module¹⁶, used for the system's visualization of mechanism diagrams, does not presently include proper support for explicit lone pair or free radical entities. Instead, the MarvinView arrows in these cases must currently be drawn as originating from an atom, despite the atom's electrons being the intended origin of these arrows.

Another nuance for these curved arrow mechanism diagrams is revealed when an electron source is a bond. In such cases, the target atom list *must* contain one of the source atoms to yield unambiguous results. In Figure 4b, this is represented by the dashed bond between atom 1 and 3, indicating a forming bond. Without this information (the dashed line or any equivalent), the reader is uncertain as to whether atom 1 or 2 should be bonded to atom 3 after the electrons have moved.

2.2.2 Stereochemistry—Many reaction examples in this manuscript have their stereochemistry simplified for clarity, but the actual system enforces that all molecules processed have complete stereochemistry specified. This ensures that any reactions that

actually have stereo specific outcomes are modeled appropriately based on fully specified inputs. To achieve this, any molecule processed by the system that contains unspecified stereocenters has all of its stereo-combinations enumerated to represent the corresponding racemic mixture (Figure 5). For unspecified E vs. Z stereochemistry of double bonds, rather than enumerate both possibilities, the system preferentially selects the isomer that keeps the largest substituents *trans* with respect to each other.

For reactions that actually have stereo specific or stereo selective outcomes, the SMIRKS language already has the expressive power to represent these transformations as illustrated in Figure 6. These representations are based on the use of the “@” symbol to specify atom chirality, and the “/” and “\” symbols to specify bond chirality¹⁴.

2.2.3 Potential Prediction Mistakes—Transformation rules as described thus far are still insufficient to provide robust reaction predictions. If alkene hydrobromination reactions were based solely on the two SMIRKS rules from Table 1, many mistakes would be made, such as those illustrated in Figure 7a, b, and c.

To develop more robust predictions that address the above issues, we must add more specific transformation rules and prioritize the list of rules by an appropriate precedence order.

2.3 Reagent-Rule Links

Reagents are basically ranked collections of elementary transformation rules. In the database implementation of the knowledge base, reagent-rule link tables assign transformation rules to reagent models, with an additional priority rank value to indicate which rules should be attempted first before descending down the precedence order. To a first order approximation, the rules are ranked in terms of the relative “reactivity” of the step being modeled. All rule links have a priority rank specified to enforce complete ordering of the rules within a reagent model, though ties are allowed to represent elementary reaction steps that are equally likely to occur (e.g., Sn1 vs. E1 termination of a carbocation intermediate). Table 3 includes a subset of the linked rules from the complete HBr reagent model. Patterns worth noting that address the issues mentioned in Figure 7 include:

1. “Carbocation, Anion Addition” is ranked higher than any “Alkene, Protic Acid Addition” to prevent production of a species with multiple positive charges.
2. Several “Alkene, Protic Acid Addition” rules exist, differentiated by what kind of carbocation each would yield. These are ranked to ensure the more stable carbocations will be formed before any others are attempted.
3. Carbocation rearrangement rules are added with high priority, and again note that several exist to account for the different possibilities, ranked respectively.

In addition to a priority ranking value, the rule link records include other supporting information and control logic. For example, warning levels and messages are attached to the carbocation rearrangement rule links in Table 3 such that, when these rules are triggered, the system can alert the user to these reaction steps that are likely to be unintended or undesirable. The system can also control the timing of certain rule triggers based on the assignment of a status number to (intermediate) reactant molecules. Rule links include a post-status number such that application of the rule not only transforms the reactant molecule structure, but also the status number associated with it. The convention established in the system is that starting material reactants begin with a status number of 0 and typically result in a status number of 100 to reflect conversion into a final product. This is particularly useful for managing reagent models representing multi-stage reactions, such as first treating a substrate with LiAlH₄ and then following with aqueous workup. To achieve this two-stage effect, the rules associated with the

LiAlH₄ step modify the input reactant's status number from 0 to 100, while the rules associated with the aqueous workup have pre-status conditions that will not allow them to trigger until the input reactants have achieved a status number of at least 100.

2.4 Reagent Models

When chemists depict reagent usage on paper, such as the hydrobromination example in Figure 2a, they typically just write “HBr” over a reaction arrow. To model what is represented there, the reagent is assigned a collection of elementary transformation rules with priority rank values, such as those in Table 3. To complete the reagent model, additional information on implied reactants and products must be tracked. This information is often neglected by chemical data systems and human chemists alike, but it is necessary to satisfy the fully balanced and atom-mapped reaction equations. For example, alkene hydrobromination reactions (Figure 2a) should be aware of implied “HBr” reactants to clearly specify the source of those product atoms. Similarly, condensation reactions (Figure 8) should be aware of implied “H₂O” products to clearly specify the fate of those reactant atoms, instead of implying that they were annihilated.

2.4.1 Recurring Reactivity Patterns and Variable Reaction Conditions—In developing the reagent models, many recurring mechanisms and reactivity patterns are found. As a result, general reaction rules can be reused in many reagent models. For example, to develop a reagent model for aqueous sulfuric acid that can add H₂O to alkenes, we can reuse almost every single rule in Table 3 without having to duplicate any effort. The only change necessary is to replace the “Carbocation, Anion Addition” rule with a similar rule for “Carbocation, Hydroxyl Addition.” This kind of rule reuse is facilitated by the reagent-rule link records which allow for a many-to-many relationship between reagents and rules. Furthermore, a simple reagent inheritance framework is supported where all of the general rules associated with a common pattern (e.g., carbocation chemistry) are linked to an abstract reagent model and then the actual, concrete reagent models that users interact with (e.g., the HBr reagent) simply “inherit” any rules from the abstract reagent model without having to reconnect each individual rule to the concrete reagent model.

This reuse of rules, representing recurring reactivity patterns, makes it convenient to develop reagent models representing similar but variable reaction conditions. For example, there are several reagent models in the system representing treatment with aqueous (sulfuric) acid, but these are qualified with variations in their reaction conditions such as “catalytic,” “cold, dilute,” “hot, dilute,” “cool, concentrated/fuming,” and “hot, concentrated.” These reagent models have identical effects in many cases, as they are based on a common set of reaction rules, but they can include additional rules or override rule priority rankings to achieve variable results. Figure 8 illustrates a reversible reaction which can be completed in either direction by selecting the appropriate reaction condition variant of an aqueous acid reagent model. These reagent models are nearly identical in reaction rule content, but the model representing a dilute acid with excess H₂O has a rule for substrate hydrolysis, which is priority ranked higher than in the complementary reagent model for concentrated acid (minimal or distilled H₂O). Finally, a few “reagent” models in the system do not even represent an actual chemical reagent, but rather fairly generic solvent conditions in cases where the reactions are driven primarily by the reactants (Table 5).

3 Results

3.1 Core Predictive Capabilities

The core reaction prediction capabilities supported by the system are illustrated in Figures 9 and 10. The products and curved arrow mechanism diagrams for these reactions are not pre-

coded results but rather are dynamically predicted within a second based on the input reactant and reagent combinations. Furthermore, the reagent models implemented in the system do not correspond to specific “name reactions¹⁷” such as a “Michael Addition” or “Claisen Condensation.” Instead, reagent models represent general reaction conditions such as treatment of the reactants with NaOH. The NaOH reagent model, for example, contains rules giving it the ability to predict many different reactivity patterns, including Sn2/E2 chemistry, nucleophilic acylations, aldol chemistry, and more. This general predictive power is possible without “memorizing” every possible “macroscopic” reaction pattern because the reagent rule sets are designed to match the “microscopic” elementary reaction steps at the mechanistic level of detailed electron flows. This allows the system to automatically derive specific overall, balanced and atom mapped, reaction patterns based on their underlying mechanisms.

Currently over 80 reagent models (listed in Table 5) have been developed for the system based on over 1,500 prioritized SMIRKS transformation rules. The examples described above were chosen for simplicity of presentation, but significantly more complex reactivity has been modeled with these rules. Reaction topics currently implemented are listed in Table 4, based on sections of content adapted from the Bruice¹⁸, Loudon¹⁹, and Smith²⁰ organic chemistry texts.

3.2 System Validation Process

The examples in Figure 9 illustrate a few specific reactions the system is known to reproduce accurately and consistently, but to ensure prediction validity across a range of possible inputs, we have manually composed over 4,500 specific reaction test cases for the system. These test cases systematically cover a range of relevant functional group combinations, including negative test cases where the most reasonable prediction is that “no reaction” will occur (e.g., treatment of a saturated hydrocarbon with acid or base). As part of a rigorous unit testing process, new test cases are added whenever the system’s rule set is expanded or modified. Before any changes to the rule set are accepted, all new and prior test cases are verified to ensure prediction validity remains intact. Furthermore, as a limited form of crowd-sourcing, users can submit a “challenge” if they believe any reaction predicted by the system is incorrect. After a few years of system service and over 3,000 users, a few dozen challenges have been submitted, but less than a handful actually identified legitimate prediction errors in the system. These few legitimate challenges alerted us to make changes in the rule set to correctly handle novel inputs, but in most of the remaining cases, the challenges came from student users who did not fully understand the chemistry involved.

3.3 Chemical Education

A specific application of the expert system that illustrates many of its predictive capabilities is a chemical education system to support the learning of organic chemistry reactions, syntheses, and mechanisms¹¹. This educational application challenges students to solve organic synthesis and mechanism elucidation problems but, unlike typical online learning applications, the underlying expert system enables teaching support for instructors and a richer learning experience for students. For instructors, this includes automated problem generation and grading. For students, this includes also automatic problem generation, as well as the fostering of inquiry-based learning²¹ where students can conduct and observe virtual experiments by selecting their own novel reactant and reagent combinations. This chemical education application has been tested in several courses at the University of California, Irvine where correlative evidence indicates that students who use the system score on average ~10% better on examinations than those who do not¹¹.

3.4 Validation of Synthesis Plans

A natural extension of the reaction prediction functionality is to apply it towards solving retro synthetic design problems²². To a large degree, all that is necessary is to take the transformation rules which normally convert reactants into predicted products and invert them to instead convert target products into proposed precursors. Making these kinds of retro synthetic suggestions is relatively straightforward and commonplace amongst computer-aided synthesis design tools^{9, 23, 24}, but the additional value gained here is that any proposed precursors can be passed back through the forward reaction prediction reagent model to validate that the intended target will actually be produced by the proposed precursors. This forward validation step can further contribute a reliability score to any proposed reaction by predicting whether a mixture of different products or stereoisomers could cripple the yield of the intended product.

3.5 Combinatorial Library Design

Another application of the reaction prediction technology is in combinatorial library design. This applies for both general library design¹⁰, where virtual molecules are systematically enumerated from an initial pool of building blocks, and targeted library design²², where virtual molecules that are structurally similar to a target compound are constructed from building blocks that are similar to substructures of the target.

The additional value of the reaction prediction technology is that it provides a natural solution to the library design problem of generating virtual compounds of reasonable synthetic feasibility²⁵. Design of a combinatorial library by enumerating all possible structures up to some constraint of atom number²⁶ or reaction types²⁷ can generate many possible structures, but leaves open the challenge of filtering down to those that could be readily synthesized. Rather than developing heuristic rules or scoring functions to estimate synthetic feasibility, systematically applying the expert system's reagent models to an initial pool of available starting materials will generate a large virtual library of compounds while simultaneously proposing a reasonable synthetic reaction to produce each compound. Furthermore, the robust reagent models allow the generation process to easily filter out any proposed reactions that would yield undesirable side reactions or mixtures as illustrated in Figure 12.

3.6 Reaction Discovery

Given the pre-programmed nature of a rule-based system, it seems unlikely that this system could discover any new types of reactions that were not already known by the knowledge engineer who authored the rules. While this may be true in terms of discovering individual elementary reaction processes, the many possible ways elementary steps can be composed into overall reactions may discover novel results. Figure 13 illustrates example reactions predicted by the system where a straightforward result was expected, but the system continued to identify and apply transformation rules for reasonable elementary reaction steps which resulted in different overall reaction patterns.

5 Discussion

A reaction expert system founded upon fundamental reaction prediction capabilities has been developed to provide a platform for addressing problems ranging from retro synthetic analysis and combinatorial library design²² to mechanism elucidation and chemical education¹¹. The prediction system is based on over 1,500 manually composed transformation rules representing fully balanced and atom-mapped elementary reaction steps with over 4,500 test cases to validate prediction accuracy and consistency.

While this rule-based approach to reaction prediction is already useful in many applications, the approach does have its limitations. Currently just over 80, most common, reagent models

are implemented in the system, but it may require hundreds of reagent model variations to achieve comprehensive coverage of the breadth of modern organic chemistry. While additional reagent models can always be added to expand the system's coverage, the size and complexity of the rule set makes progressive addition of rules increasingly more challenging.

An alternative approach to manually composing reaction rules is to automatically generate the rules by perceiving common patterns from reaction databases^{28, 29}. The automation of these approaches is certainly appealing, but the depth of prediction models they can generate is often limited by the reaction databases available for them to work from. Full access to large reaction databases is often highly restricted and, even if data mining access is possible, the data is noisy and tends to lack fully balanced and atom-mapped reactions. Furthermore, these data almost always describe overall "macroscopic" reactions with no detail on the underlying mechanisms and elementary reaction steps.

To achieve a greater level of generality and robustness, an alternative reaction predictor design could instead be driven by more fundamental principles of molecular orbital theory³⁰ and reaction kinetics simulations, though it would probably do so at the cost of longer prediction times. Even for the development of such a principle-driven approach, the rule-based system described here can be useful for generating a virtual database of detailed reaction mechanisms with fully balanced and atom-mapped reaction data to train and validate the principle-driven system. In the meantime, the rule-based expert system already provides a platform for solving a range of important chemistry applications by addressing the fundamental problem of chemical reaction prediction, at a unique mechanistic level of detail, with sub-second prediction times.

Acknowledgments

Work supported by an NIH Biomedical Informatics Training grant (LM-07443-01) and NSF grants EIA-0321390 and 0513376 to PB. We acknowledge OpenEye Scientific Software, Peter Ertl of Novartis (JME Editor), and ChemAxon for academic software licenses. We thank Drs. Suzanne Blum, Zhibin Guan, Elizabeth Jarvo, Susan King, Larry Overman, Scott Rychnovsky, Kenneth Shea, Mare Taagepera, David Van Vranken, Chris Vanderwal and Gregory Weiss and their students for their feedback and usage of the system in their chemistry courses. We acknowledge Matthew A. Kayala, Peter Phung, and Paul Rigor for contributing to software design and development. We thank Dr. James Nowick for additional feedback and comments.

References

1. Jorgensen WL, Laird ER, Gushurst AJ, Fleischer JM, Gothe SA, Helson HE, Paderes GD, Sinclair S. CAMEO: a program from the logical prediction of the products of organic reactions. *Pure Appl Chem* 1990;62(10):1921–1932.
2. Hollering R, Gasteiger J, Steinhauer L, Schulz KP, Herwig A. Simulation of Organic Reactions: From the Degradation of Chemicals to Combinatorial Synthesis. *J Chem Inf Comput Sci* 2000;40:482–494. [PubMed: 10761155]
3. Sello G. Reaction prediction: the suggestions of the Beppe program. *J Chem Inf Comput Sci* 1992;32(6):713–717.
4. Socorro IM, Taylor K, Goodman JM. ROBIA: A Reaction Prediction Program. *Org Lett* 2005;7(16):3541–3544. [PubMed: 16048337]
5. Satoh H, Itono S, Funatsu K, Takano K, Nakata T. A Novel Method for Characterization of Three-Dimensional Reaction Fields Based on Electrostatic and Steric Interactions toward the Goal of Quantitative Analysis and Understanding of Organic Reactions. *J Chem Inf Comput Sci* 1999;39(4):671–678.
6. Benko G, Flamm C, Stadler PF. A Graph-Based Toy Model of Chemistry. *J Chem Inf Comput Sci* 2003;43(4):1085–1093. [PubMed: 12870897]
7. Gasteiger, J.; Engel, T. *Chemoinformatics: A Textbook*. Wiley-VCH; 2003.
8. Hemmer, MC. *Expert Systems in Chemistry Research*. CRC Press; Boca Raton, FL: 2008.

9. Todd MH. Computer-Aided Organic Synthesis. *Chem Soc Rev* 2004;34(3):247–266. [PubMed: 15726161]
10. Schnur DM. Recent trends in library design: ‘rational design’ revisited. *Curr Opin Drug Discovery Dev* 2008;11(3):375–380.
11. Chen JH, Baldi P. Synthesis Explorer: A Chemical Reaction Tutorial System for Organic Synthesis Design and Mechanism Prediction. *J Chem Educ* 2008;2008(85):1699.
12. Miller, A. Writing Reaction Mechanisms in Organic Chemistry. Academic Press; San Diego, CA: 1992.
13. Grossman, R. The Art of Writing Reasonable Organic Reaction Mechanisms. Vol. 2. Springer; New York, NY: 2003.
14. James, CA.; Weininger, D.; Delany, J. Daylight Theory Manual. [accessed Apr 2009]. <http://www.daylight.com/dayhtml/doc/theory/theory.toc.html>
15. OpenEye . OEChem Toolkit. 2009. 1.4.2; <http://www.eyesopen.com>
16. ChemAxon . Marvin Applets. 2009. 5.1.2; <http://www.chemaxon.com>
17. Li, JJ. Named Reactions: A Collection of Detailed Reaction Mechanisms. Vol. 3. Springer; Berlin, Germany: 2006.
18. Bruice, PY. Organic Chemistry. Vol. 4. Prentice-Hall; Upper Saddle River, NJ: 2004.
19. Loudon, M. Organic Chemistry. Vol. 4. Oxford University Press; 2001.
20. Smith, JG. Organic Chemistry. Vol. 2. McGraw-Hill; Boston, MA: 2006.
21. Joolingen, WRv; Jong, Td; Dimitrakopoulou, A. Issues in computer supported inquiry learning in science. *J Comp Asst Learn* 2007;23:111–119.
22. Chen JH, Linstead E, Swamidass SJ, Wang D, Baldi P. ChemDB Update - Full- Text Search and Virtual Chemical Space. *Bioinformatics* 2007;23(17):2348–2351. [PubMed: 17599932]
23. Corey EJ, Long AK, Rubenstein SD. Computer-Assisted Analysis in Organic Synthesis. *Science* 1985;228(4698):408–418. [PubMed: 3838594]
24. Hanessian S. Man, machine and visual imagery in strategic synthesis planning: Computer-perceived precursors for drug candidates. *Curr Opin Drug Discovery Dev* 2005;8(6):798–819.
25. Allu TK, Oprea TI. Rapid Evaluation of Synthetic and Molecular Complexity for in Silico Chemistry. *J Chem Inf Model* 2005;45(5):1237–1243. [PubMed: 16180900]
26. Fink T, Reymond J-L. Virtual Exploration of the Chemical Universe up to 11 Atoms of C, N, O, F: Assembly of 26.4 Million Structures (110.9 Million Stereoisomers) and Analysis for New Ring Systems, Stereochemistry, Physicochemical Properties, Compound Classes, and Drug Discovery. *J Chem Inf Model* 2007;47(2):342–353. [PubMed: 17260980]
27. Cramer RD, Soltanshahi F, Jilek R, Campbell B. AllChem: generating and searching 10²⁰ synthetically accessible structures. *J Comput-Aided Mol Des* 2007;2007(21):341–350. [PubMed: 17253118]
28. Chen L, Gasteiger J, Rose JR. Automatic Extraction of Chemical Knowledge from Organic Reaction Data: Addition of Carbon-Hydrogen Bonds to Carbon-Carbon Double Bonds. *J Org Chem* 1995;60(24):8002–8014.
29. Law J, Zsoldos Z, Simon A, Reid D, Liu Y, Khew SY, Johnson AP, Major S, Wade RA, Ando HY. Route Designer: A Retrosynthetic Analysis Tool Utilizing Automated Retrosynthetic Rule Generation. *J Chem Inf Model* 2009;49(3):593–602. [PubMed: 19434897]
30. Fleming, I. Frontier Orbitals and Organic Chemical Reactions. Wiley; New York, NY: 1976.

Figure 1.

Retro synthetic suggestions, illustrating the need for reaction validation capabilities. The first example illustrates a simple benzyl alcohol target compound and a proposed pair of precursor molecules to synthesize the target by a Grignard reaction. The second example illustrates a nearly identical target compound and the precursors that would be proposed by naively applying the analogous retro synthetic transformation. This second suggestion is invalid because it does not consider the acid-base, side reaction between the alcohol and organometallic reagent that will ruin the intended result.

Figure 2.

Figure 2a. Representation for the overall “macroscopic” reaction of an alkene with hydrobromic acid, indicating the starting material, reagent, and final product. In the context of the system, the alkene starting material reactant and the selection of “HBr” as a reagent represents the expected input, while the alkyl bromide product represents the primary output. Figure 2b. Detailed reaction mechanism for an alkene hydrobromination reaction, illustrating the underlying “microscopic” elementary processes that the overall reaction is based upon. This represents the detailed expected output when applying a reagent model for hydrobromic acid to the alkene reactant.

Figure 3.

Overall architecture of the system. The knowledge base is implemented in a database and the right column provides a simplified view of the database schema. There exists a one-to-many relationship between reagents and reagent-rule links and likewise between transformation rules and reagent-rule links. The combination of the previous relationships creates a many-to-many relationship between reagents and rules.

Figure 4.

Figure 4a. Arrow pushing mechanism diagram generated when applying the SMIRKS reaction transformation rule [C+:1].[-:2]>>[C+0:1][+0:2] and electron flow specification 2=1 to a carbocation electrophile and a bromide anion nucleophile. This represents the movement of two electrons from atom 2 to atom 1.

Figure 4b. Arrow pushing mechanism diagram generated when applying the SMIRKS reaction transformation rule C:1=[C:2][c:10].[H:3][Cl,Br,I:4]>>[H:3][C:1][C+:2][c:10].[Cl,Br,I:0:4] and electron flow specification 2,1=1,3;3,4=4 to styrene in hydrobromic acid, representing protonation of the nucleophilic π orbital electrons. The dashed line between carbon 1 and hydrogen 3 indicates the site of a forming bond in the mechanism step.

Figure 5.

Products generated by Sn1 and E1 reactions applied to an achiral benzyl carbocation. No stereo selectivity exists for this Sn1 substitution reaction, so the system will enumerate both enantiomers as possible products of the reaction to represent a racemic mixture. The E1 elimination reaction could theoretically produce two diastereomeric products, but the system will preferentially select the one which keeps the largest substituents *trans* with respect to each other, to reflect the typical pattern of stereo selectivity in these reactions.

```
[*:10]/[CX3;$(*[O,N,S]):1]=[C:2]-[C:3]=[CX3:4]/[*:11].
[*:12]\[CX3:5]=[CX3:6]/[$(*[O,N,S]),$(C#N):13]>>
[*:10][CX4;@:1]1-[C:2]=[C:3]-[CX4;@@:4]([*:11])[CX4;@:5]([*:12])[CX4;@@:6]1[*:13]
```

Figure 6.

Product prediction for a Diels-Alder reaction using the accompanying SMIRKS transformation rule. The example illustrates the expressive power of the rule to enforce the regioselectivity, stereospecificity, and stereoselectivity of the reaction. Regioselectivity: Carbon 1 preferentially assumes an *ortho* position with respect to carbon 6, based on the pattern of their substituents. Stereospecificity: The (*E,E*) diene results in substituents at carbon 1 and 4 that are *syn* with respect to each other, and likewise for the *Z* dienophile resulting in *syn* substituents at carbon 5 and 6. Stereoselectivity: Assuming kinetic preference for the *endo* product over the *exo* product, the orientation of stereocenters at carbon 1 and 4 is defined with respect to those at carbon 5 and 6. This example illustrates one product generated by the stereospecific Diels-Alder reaction, but the actual reaction will yield a racemic mixture including the enantiomeric product. In such cases, a second copy of the SMIRKS rule is included. The second rule is modified with inverted stereospecification symbols to yield the respective enantiomeric product.

Figure 7.

Figure 7a. Potential prediction mistake: Adding another proton to a carbocation intermediate, yielding an unlikely species containing multiple atoms with a positive charge.

Figure 7b. Potential prediction mistake: Ignoring regioselective preference for carbocations on more substituted carbons. Protonation that yields a secondary carbocation is highly unlikely when the protonation could occur on the other end of the double bond, yielding a tertiary carbocation of greater stability.

Figure 7c. Potential prediction mistake: Ignoring the possibility of unintended reactivity such as carbocation rearrangements. The anion (Br^-) does not add directly to the secondary carbocation. Instead, because the carbocation is adjacent to a tertiary center, a 1,2 hydride shift will alter the site of the carbocation.

Figure 8.

Reversible Fischer esterification reaction whose preferred outcome can be altered by selecting a reagent model representing a particular set of reaction conditions. In this example, both the forward and reverse reaction can be driven using an aqueous acid reagent, but selection of a concentrated acid, under heat, will tend to distill any H_2O , driving the reaction towards the ester. Selection of a dilute acid with excess H_2O solvent will drive the reaction towards the hydrolyzed carboxylic acid and alcohol. The reagent models for these two cases are nearly identical in their reaction rule content, but differences are achieved by respectively adjusting the priority rank of the hydrolysis vs. esterification rules in each reagent model.

Figure 9.

Progressively more complex reactions predicted by the system (*Sn*₂ substitution, nucleophilic acylation/saponification, Robinson annulation), all based on a single common reagent model (NaOH). This reagent model contains relevant reactivity rules to represent a set of general reaction conditions as opposed to one rule for each specific reaction.

Figure 10.

Reaction mechanism details page generated by the system to illustrate the chain of elementary reaction steps used to predict the outcome of the Robinson annulation reaction at the end of Figure 9. Each step includes a system-generated curved arrow mechanism diagram and an accompanying verbal description. Some steps include additional informative or cautionary notes to assist the user.

Figure 11.

Figure 11a. Examples of major pharmaceutical drugs and synthetic reactions proposed by naively applying a typical retro synthetic pattern matching approach. The expert system's robust reagent models can provide the critical "expertise" of synthesis plan validation by identifying all of these proposed plans as ineffective due to unintended side reactions.

Fenofibrate: The proposed organostannane precursor is difficult to prepare in the presence of another aryl halide. Bupropion: Over-oxidation of the benzylic alcohol is likely.

Albuterol: Organometallic Grignard reagent cannot be produced in the presence of acidic OH groups.

Figure 11b. Examples of possible synthetic reactions for several pharmaceutical drugs that the expert system's reagent models can reproduce, validating the intended reaction products.

Figure 12.

Flowchart for designing a combinatorial library using the expert system's reagent models. A collection of available building blocks is passed through the reagent models in all (pair-wise) combinations to predict reasonable virtual products along with a specific synthetic reaction proposal for each. The reagent models will naturally sort the results into relevant subsets based on the proposed synthetic reactions. An example product and respective synthetic reaction is illustrated for each of these subsets. Acceptable: Products generated from proposed reactions that the system validates as reasonable and effective for use in the library. Unreliable: Products generated from proposed reactions that may work, but are likely to produce an unreliable mixture of many side products. Invalid: Products for which no proposed synthetic reaction is acceptable, due to side reactions that will disrupt the intended result.

Figure 13.

Reactions predicted by the system with results that defied straightforward expectations. The results are based on additional combinations of elementary reaction steps that the system recognizes can be chained together. The first reaction is expected to be a simple hydride reduction of the imine to produce a nitrogen anion that is subsequently neutralized by the protic solvent. Instead, the system recognizes that the nitrogen anion intermediate is a strong nucleophile that can attack the nearby ester to form a lactam before the solvent neutralization. The second reaction is expected to open the epoxide by the enolate nucleophile. The system does predict that effect, but it also recognizes the epoxide opens up to yield an oxygen anion which is a strong nucleophile that can reach back to attack the original ester to form a lactone.

Table 1

SMIRKS transformation rules corresponding to a simple alkene hydrobromination reaction model. Each item in brackets corresponds to an atom in the reaction equation. The “>>” symbol delimits reactants from products. The numbers following colons are atom-map indexes used to specify which reactant atoms correspond to which product atoms. Further specification of the SMIRKS language can be found in the references¹⁴.

SMIRKS	Description
[C:1]=[C:2].[H:3][Cl,Br,I,\$(OS=O):4]>>[H:3][C:1][C+:2].[-:4]	Alkene, Protic Acid Addition
[C+:1].[-:2]>> [C+0:1][+0:2]	Carbocation, Anion Addition

Table 2

Definition of the symbols that can be used in the electron flow specification language.

Symbol	Description
;	Delimits specifications for diagrams with multiple electron flow arrows.
=	Represents a double-headed arrow for the movement of 2 electrons, delimiting source from sink atoms.
-	Represents a single-headed “fishhook” arrow for the movement of 1 electron (i.e., free radical reactions), delimiting source from sink atoms.
n_i	Numerical indexes which identify atoms representing sources and sinks for the arrows.
,	Atom delimiter for when the source or target of the arrow consists of multiple atoms (i.e., 2 atoms specified to indicate bond electrons). The order in which the atoms are listed here does not affect the resulting diagram.

Table 3

Example of 10 prioritized transformation rules, relating to alkene hydrobromination reactions, out of the 92 rules used in the complete robust HBr reagent model. Included for each transformation rule is not only the SMIRKS pattern and description, but a relative priority rank to indicate the order in which the rules should be attempted. The existence of several variants for similar rules and the customized priority ordering enables robust reaction predictions that address the issues noted in Figure 7. An electron flow specification accompanies each rule to support curved arrow mechanism diagrams.

SMIRKS	Description	Electron Flow	Priority
[H:10][CH1:1][C+;!H0:2]>>[C+:1][C+0:2][H:10]	Carbocation, Hydride Shift from Tertiary	1,10=10,2	10
[C:10][CH0:1][C+;!H0:2]>>[C+:1][C+0:2][C:10]	Carbocation, Methyl Shift from Quaternary	1,10=10,2	9
[H:10][CH2:1][CH2+:2]>>[C+:1][C+0:2][H:10]	Carbocation, Hydride Shift from Secondary	1,10=10,2	8
[C+:1].[-:2]>>[C+0:1][+0:2]	Carbocation, Anion Addition	2=1	7
[C:1]=[C;\$(*O):2].[H:3][Cl,Br,I,\$(OS=O):4]>>[H:3][C:1][C+:2].[-:4]	Alkene, Protic Acid Addition, Alkoxy	2,1=1,3;3,4=4	6
[C:1]=[C;\$(*a):2].[H:3][Cl,Br,I,\$(OS=O):4]>>[H:3][C:1][C+:2].[-:4]	Alkene, Protic Acid Addition, Benzyl	2,1=1,3;3,4=4	5
[C:1]=[C;\$(*=):2].[H:3][Cl,Br,I,\$(OS=O):4]>>[H:3][C:1][C+:2].[-:4]	Alkene, Protic Acid Addition, Allyl	2,1=1,3;3,4=4	4
[C:1]=[CH0:2].[H:3][Cl,Br,I,\$(OS=O):4]>>[H:3][C:1][C+:2].[-:4]	Alkene, Protic Acid Addition, Tertiary	2,1=1,3;3,4=4	3
[C:1]=[CH1:2].[H:3][Cl,Br,I,\$(OS=O):4]>>[H:3][C:1][C+:2].[-:4]	Alkene, Protic Acid Addition, Secondary	2,1=1,3;3,4=4	2
[C:1]=[C:2].[H:3][Cl,Br,I,\$(OS=O):4]>>[H:3][C:1][C+:2].[-:4]	Alkene, Protic Acid Addition	2,1=1,3;3,4=4	1

Table 4

List of reaction topics currently covered in the system. Section numbers correspond to the Loudon textbook, though the system is not tied to any particular content source since it is designed to model the fundamental underlying chemistry. Gaps in the section numbers correspond to textbook chapters that do not include any relevant reactions to model, such as chapters on stereochemistry or spectroscopy.

Section	Description
5	Alkenes
9.04	Substitution (Nucleophilic) of Alkyl Halides
9.05	Elimination Reactions of Alkyl Halides
10	Alcohols and Epoxides
11.04	Epoxides and Organometallic Compounds
11.05	Oxidation of Alcohols and Alkenes
14	Alkynes
15	Dienes, Conjugation, Diels-Alder
16	Electrophilic Aromatic Substitution
17	Allylic and Benzylic Reactivity
17.02	Alkanes, Radical Reactions
18	Transition Metal (Pd) Catalysis
18.04	SnAr and Benzyne Reactions
19	Aldehydes and Ketones
20.1	Redox of Alcohols and Carbonyls
21	Carboxylic Acid Derivatives
22	Enolate Chemistry
22.04	Aldol Chemistry and Michael Addition
22.05	Claisen Condensations
22.08	Organometallic Addition, Conjugate Addition
23	Amines
23.1	Arenediazonium Reactions
24	Naphthalene and Heteroaromatic EAS Reactions
24.05	Pyridine Derivatives
25	Pericyclic Reactions
26.04	Amino Acid Synthesis
26.07	Peptide Synthesis
27	Carbohydrates

Table 5

Listing of 80 reagent models currently implemented in the system that users can combine with reactant molecules to predict the course and major products of reactions. A few models do not reflect an actual chemical reagent, but instead represent reactions driven primarily by the reactants in a generic solvent. In particular, there are “reagents” for simply mixing the reactants in different solvent types (“Mix Reactants, Protic” and “Mix Reactants, Aprotic”) and one for mixing the reactants under heat to model thermal pericyclic reactions (“Pericyclic Reactions (thermal)”).

Reagent Model Descriptions	
Pd(0) (catalyst)	Clemmensen Reduction (acid)
Pericyclic Reactions (thermal)	Wolff-Kishner Reduction (base)
Mix Reactants, Aprotic	Oxidation (base, permanganate)
Mix Reactants, Protic	Oxidation (acid, chromate)
Hydrogen Fluoride (Friedel-Crafts Catalyst)	Oxidation (MnO ₂ , benzylic, partial)
Lewis Acid (Friedel-Crafts Catalyst)	Oxidation (PCC)
Sulfuric Acid (catalytic)	Oxidation (Nitric Acid)
Sulfuric Acid (cold, dilute)	SOCl ₂
Sulfuric Acid (hot, dilute)	PBr ₃
Sulfuric Acid (cool, concentrated/fuming)	Tosylation
Phosphoric Acid (hot, concentrated)	Triflate Preparation
Bromination, Lewis Acid	POCl ₃
Nitric Acid	P ₂ O ₅
NaOH	Acetic Anhydride
NaOEt	DCC
NaH	Mg (Grignard)
NaNH ₂	Lithium
LDA	Organocuprate Preparation
Hydroboration-Oxidation	Organostannane Preparation
Hydrobromination	Bromination, Radical
Hydrobromination, Peroxide	NBS, Peroxide
Bromination	PPh ₃ , BuLi (Phosphonium Ylide Prep)
Bromohydrin	TMSCl, Et ₃ N
Br ₂ , H ₃ O ⁺	Fmoc Amine Protection
Br ₂ , NaOH	NH ₄ ⁺ F ⁻
Hydrogenation	Fmoc Deprotection (Piperidine)
Hydrogenation, Partial	TFA (para-oxy benzyl deprotection)
Hydrogenation, Pd/BaSO ₄	NH ₄ Cl, NaCN
Na, NH ₃	Cyanohydrin
O ₃ , CH ₃ SCH ₃	Benzylic Halide (para-oxy)
O ₃ , H ₂ O	Arenediazonium Prep (HCl)
OsO ₄ , NaHSO ₃ (syn dihydroxylation)	Arenediazonium Prep (H ₂ SO ₄)
Periodic Acid (HIO ₄)	Arenediazonium Prep (HBr)
Peroxyacid (mCPBA)	Arenediazonium (F)
Sharpless Epoxidation (+)-DET	Arenediazonium (Cl)
Sharpless Epoxidation (–)-DET	Arenediazonium (Br)

Reagent Model Descriptions	
LiAlH ₄	Arenediazonium (I)
DIBALH	Arenediazonium (C#N)
NaBH ₄	Arenediazonium (H) Hypophosphorus Acid
NaBH ₃ CN	Hofmann Elimination