

Algorithmic Analysis of Cahn–Ingold–Prelog Rules of Stereochemistry: Proposals for Revised Rules and a Guide for Machine Implementation

Robert M. Hanson*¹ and Sophia Musacchio²

Department of Chemistry, St. Olaf College, 1520 St. Olaf Avenue, Northfield, Minnesota 55057, United States

John W. Mayfield³

NextMove Software Ltd, Unit 23, Cambridge Science Park, Cambridge, CB4 0EY, United Kingdom

Mikko J. Vainio⁴


Varian Medical Systems Finland Oy, Paciuksenkatu 21, Helsinki 00270, Finland

Andrey Yerin⁵ and Dmitry Redkin⁶

Moscow Department, Advanced Chemistry Development, 6 Akademika Bakuleva Street, Moscow 117513, Russia

S Supporting Information

ABSTRACT: The most recent version of the Cahn–Ingold–Prelog rules for the determination of stereodescriptors as described in *Nomenclature of Organic Chemistry: IUPAC Recommendations and Preferred Names 2013* (the “Blue Book”; Favre and Powell. Royal Society of Chemistry, 2014; <http://dx.doi.org/10.1039/9781849733069>) were analyzed by an international team of cheminformatics software developers. Algorithms for machine implementation were designed, tested, and cross-validated. Deficiencies in Sequence Rules 1b and 2 were found, and proposed language for their modification is presented. A concise definition of an additional rule (“Rule 6”, below) is proposed, which succinctly covers several cases only tangentially mentioned in the 2013 recommendations. Each rule is discussed from the perspective of machine implementation. The four resultant implementations are supported by a 300-compound validation suite in both 2D and 3D structure data file (SDF) format as well as SMILES (<https://cipvalidationsuite.github.io/ValidationSuite>). The validation suites include all significant examples in Chapter 9 of the Blue Book, as well as several additional structures that highlight more complex aspects of the rules not addressed or not clearly analyzed in that work. These additional structures support a case for the need for modifications to the Sequence Rules.



■ INTRODUCTION

In the 60+ years since the introduction of Cahn–Ingold–Prelog Sequence Rules in 1956,¹ the “CIP Rules” have become an integral part of chemical nomenclature, providing a way to identify the spatial arrangement of atoms of a molecule using simple mostly atom- or bond-based stereodescriptors. Over the course of this time, various authors have pointed out deficiencies in the rules and proposed solutions in the form of modifications and subrules,^{2,3} to the point where today we have eight distinct Sequence Rules: 1a, 1b, 2, 3, 4a, 4b, 4c, and 5.

In order to provide a single resource summarizing the state of the evolving rules, the International Union of Pure and Applied Chemistry (IUPAC) published the first comprehensive description of the CIP Rules in *Nomenclature of Organic Chemistry: IUPAC Recommendations and Preferred Names 2013* (referred to below as “BB 2013”).⁴ This description, an

impressive 197-page chapter with more than two hundred examples, presents a complete set of CIP Rules, along with detailed procedures for their application.

The “CIP descriptors” generated by these rules are primarily intended for use in chemical nomenclature. It would be possible to use them in other applications, such as the removal of redundant stereo specifications, the determination of structure equivalence, or canonical labeling. However, CIP-based algorithms may be more resource consuming and complex than simpler, more efficient cheminformatics algorithms developed specifically for those purposes.

We note that discoveries of deficiencies in the original rules have been made before in the context of developing computer-based implementations. This was certainly the case in 1993,

Received: May 25, 2018

Published: July 30, 2018

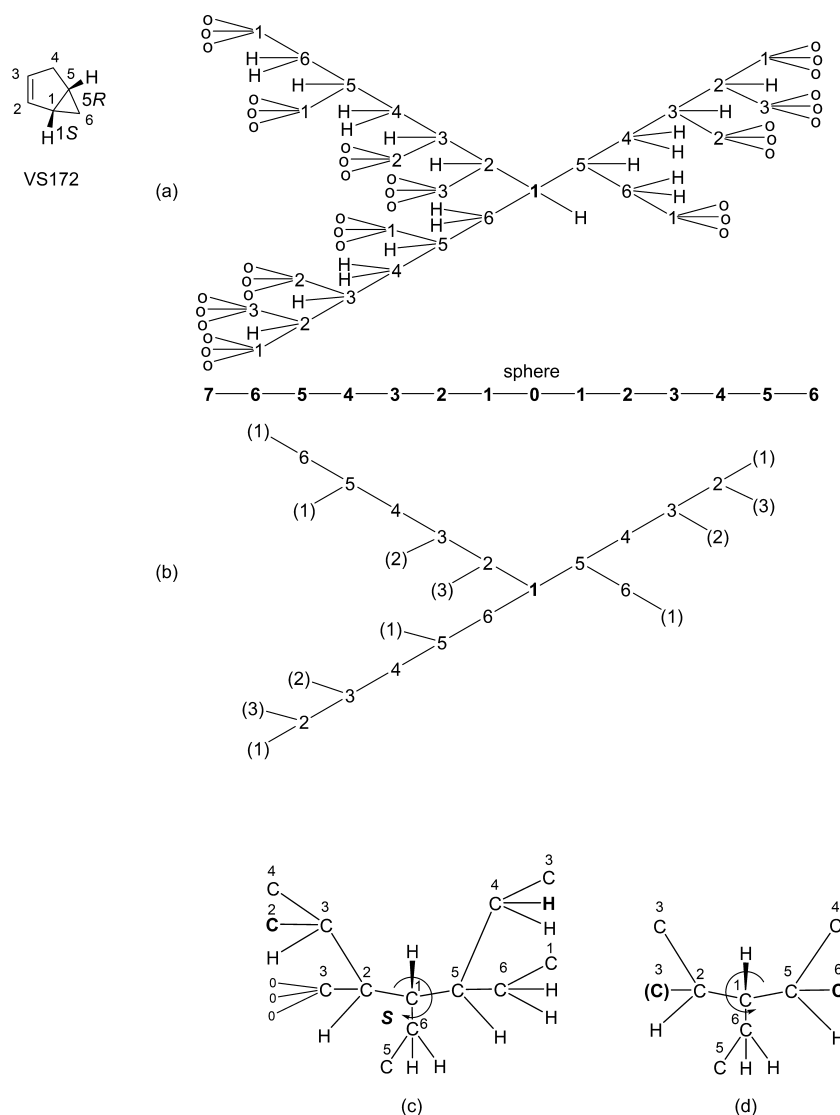


Figure 1. Digraph examples for (1*S*,5*R*)-bicyclo[3.1.0]hex-2-ene (VS172). (a) Full digraph showing carbon numbers, with hydrogens and phantom atoms (o). (b) Simplified digraph showing only carbons. (c) Partial digraph showing the proper analysis leading to an *S* descriptor for C-1 with deciding nodes in Sphere 3 in bold. (d) Improper analysis based solely on a simplified digraph that supposes that a carbon atom will always have higher priority than a duplicate carbon node.

when subrules 4a, 4b, and 4c were proposed by a group developing a stereochemical module for the LHASA computer-aided synthesis analysis program.^{5,6}

Not surprisingly, predating the open-source collaborative environment and not having a concise reference in hand, machine implementations of the CIP Rules to date have been only marginally successful. Certainly, many software developers have implemented the CIP Rules to one extent or another, but recent analysis of available software packages clearly demonstrates that there is much disagreement among these implementations, even for relatively simple compounds, among several highly respected software packages.⁷

In the spring of 2017, concurrent discussions started in two forums, the IUPAC Blue Book Project,⁸ focused on preparation of errata for BB 2013, and the Blue Obelisk Group,⁹ focused on implementation issues of CIP rules in algorithmic form. The result of these lively discussions was a coordinated and thorough analysis of the CIP Rules, involving concurrent development and improvement of four software packages: Jmol,¹⁰ Centres,¹¹ ChemSketch,¹² and Balloon.^{13,14}

In the analysis that follows, structure numbers starting with “VS” refer to compounds in the validation suite, which is available as Supporting Information and at <https://cipvalidationsuite.github.io/ValidationSuite>.

Our joint efforts have ensured that multiple, validated, independent machine implementations of the CIP Rules are made available to the cheminformatics community, as well as to anyone interested in stereochemistry and chemical nomenclature. In fact, the results of our work have already been incorporated into an interactive web site utilizing JSmol.¹⁵

Preliminary Considerations: Digraphs. As we sought consensus in our discussions, we turned to the use of “hierarchical digraphs”³ (Figure 1) constructed as described in BB 2013 P-92.1.4 for each stereogenic unit of interest. These graphs represent the chemical structure as a *finite acyclic directed graph*—a set of “nodes” and associated “edges” that radiate out from a central “root” vertex in a series of “spheres” and involve no cyclic pathways. The root vertex in each case represents the specific stereogenic unit for which we desire to

determine a stereochemical descriptor. The use of digraphs presumes a localized-bond model. Double bonds introduce pairs of terminal “duplicate nodes” that are the same nature as their “duplicated atoms” but are connected to three “phantom atoms” that have zero atomic number. Triple bonds introduce two such pairs. Each ring-closure bond introduces a duplicate node for the revisited (“ring closure”) atom. Phantom atoms and hydrogen atoms are often not shown on “simplified” digraphs (Figure 1b and d). While our implementations generate the internal machine equivalent of full digraphs (Figure 1a and c), discussions generally involve simplified versions.

While it is common to find simplified digraphs that do not show hydrogen or phantom atoms, only fully elaborated digraphs should be used for assignment of stereodescriptors. A critical aspect of the algorithm must be finding the *first* difference in two ligands, and this may not be obvious from a simplified digraph. For example, there is no rule that “real atoms have higher priority than duplicate nodes”. This is intentional. While generally true, this statement hides the fact that duplicate nodes lose to their real counterparts only in the *next higher sphere*, where their associated phantom atom, with atomic number zero, always loses to any real atom. Figure 1d shows how naïve application of the pseudorule “real atoms have higher priority than duplicate nodes” leads to a 1R descriptor rather than correct 1S assigned in accord with Rule 1a.

The use of digraphs in the determination of stereodescriptors has been important for many years. They were introduced by Prelog and Helmchen in 1982 in relation to cyclic structures³ and emphasized later by Mata et al.,⁵ in relation to the process of developing Rule 4:

If, for a clear analysis of the ligands, they must be converted into hierarchical digraphs, then the comparison must always be done considering the hierarchical digraph, not the real ligand.

Our consensus interpretation of this statement is that, at least for cyclic compounds, a unique digraph must be generated for each center in question. In practice, all of our implementations generate a machine equivalent of a digraph for each stereogenic unit in all cases, cyclic or not.

From the onset, we agreed that analysis of digraphs would always be the final arbiter in case of disputes. In such events, we allowed for three possible conclusions: (a) that the disagreement was due to different interpretations of CIP rules among software developers, (b) that there was a problem with an algorithm or its implementation in code, or (c) that the CIP rules themselves were flawed. In fact, we encountered all three of these possibilities in the process of coming to consensus, including the discovery of a small number of errors in the Blue Book, two minor flaws in the CIP rules, and a proposal for a new rule.

Preliminary Considerations: Auxiliary Descriptors. Temporary “auxiliary descriptors” are assigned solely on the basis of a given digraph for a particular stereogenic unit in question and may or may not be the “final” descriptors ultimately used to describe those centers in the end. Figure 2 shows an example (VS279) where only a minority of the auxiliary descriptors are the same as the final descriptors for the corresponding atoms.

Generation of a complete digraph, including all auxiliary descriptors, including *seqcis* (“Z”) and *seqtrans* (“E”), is required prior to application of Rule 3. Note that in the determination of auxiliary descriptors, the entire sequence of




Figure 2. Compound (VS279) with a stylized digraph for the determination of the *S* stereodescriptor on the far-left ring carbon. Auxiliary descriptors are shown with a subscript “o”. Only four of the ten auxiliary descriptors on the digraph are their final values.

all rules must be carried out for each auxiliary center on the *same* digraph, despite the fact that for the unit in question only Rules 1a, 1b, and 2 have been completed at this point. Generation of auxiliary descriptors must start from the highest sphere, proceeding toward the root. In this way, all auxiliary descriptors in higher spheres than the one being determined are already assigned. This is sufficient, as the descriptor for an auxiliary center does not depend upon any descriptor between it and the root, since the priority of a ligand leading back to the digraph root will always be ranked by Rule 1a, with no need to consider auxiliary centers. This postulate follows from the fact that auxiliary centers are always offset from the root of a digraph, and so the path back to the root is always unique in connectivity and atomic numbers.

Sequence Rules and Subrules. While BB 2013 in P-92.1.3 specifies that each of Rules 1–5 “must be exhaustively applied in the order given”, it may not be clear what exactly “exhaustively” means or how this relates to subrules such as 1a and 1b. In fact, each subrule must be applied sequentially. Effectively, there are eight (nine, if one adds our proposed Rule 6) fully independent explicit Sequence Rules: 1a, 1b, 2, 3, 4a, 4b, 4c, and 5. The fact that these are not “Sequence Rules 1–8” is simply a result of the historical evolution of the Sequence Rules. “Exhaustively” simply means “until a decision is reached, or it is determined that no such decision is possible”.

Ranking and Comparing Ligands. As described clearly in BB 2013, application of all eight rules involves the same two processes: Ligands are ranked sphere by sphere, branch by branch in a breadth-first fashion. Then, in pairs, two ligands are compared atom by atom, in order of that ranking. In practice, it is not always necessary to completely rank a ligand. Ranking—at least through Rule 2—can be carried out on a “need-to-know” basis, skipping whole sub-branches of the digraph where a decision has already been made.

RESULTS AND DISCUSSION—ALGORITHMIC ANALYSIS OF CIP PROCEDURES

What follows is a discussion of each of the eight independent Sequence Rules from the perspective of machine implementation, along with a proposal for a ninth rule that we are calling “Rule 6”.

Rule 1a: *Higher atomic number precedes lower.*

On the face of it, Rule 1a sounds simple enough to implement. And it is, except for the special cases discussed in BB 2013 P-92.1.4.4 in relation to compounds and ions with multiple chemically equivalent Kekulé structures, such as benzene, pyridine, or cyclopentadienyl anion. That section briefly introduces the idea of “atomic number averaging” for mancude-ring systems, which are “rings having (formally) the maximum number of noncumulative double bonds, e.g.

benzene, indene, indole, 4*H*-1,3-dioxine".¹⁶ The idea is to average the atomic number of the duplicate node when it is involved in multiple resonance structures. Our reading of Section P-92.1.4.4 is that it is not a well-crafted guideline, with many unanswered questions. What exactly defines the pertinent cases for which this rule applies? What about acyclic cases such as allyl anion or acetoacetate? How exactly is the averaging to be done—Do we need to know the exact weighting of all the possible resonance structures, or is it sufficient to average over just the adjacent atoms involved in the electron delocalization?

Without going into detail here, suffice it to say that atomic number averaging is a difficult procedure to describe or implement, other than the case of all-carbon neutral species, for which it is unnecessary, and some simple heterocyclic systems, such as pyridine derivatives. Thus, consider substituted pyridine VS032/VS033, Figure 3. Notice that




Figure 3. Without atomic number averaging in Rule 1a, the two chemically equivalent Kekulé structures give different assignments at the stereocenter. For VS032, the decision compares {N(7), N(7), C(6)} to {N(7), N(7), H(1)}, whereas for VS033, the decision is {N(7), C(6), C(6)} vs {N(7), N(7), H(1)}. Atomic number averaging removes the issue, comparing atomic numbers {N(7), N(7), C(13/2), C(6)} to {N(7), N(7), H(1)} in both cases.

without this consideration, Rule 1a gives two different results in this case, depending upon the choice of Kekulé structure. The correct result, *S*, derives from assigning the duplicate node for the aromatic nitrogen an averaged atomic number of 6.5, which loses to the unaveraged atomic-number 7 duplicate node of the imine nitrogen.

Rule 1b: A duplicate atom node whose corresponding nonduplicated atom node is the root or is closer to the root ranks higher than a duplicate atom node whose corresponding nonduplicated atom node is farther from the root.

Shortly into our study it became clear that the current IUPAC recommendation for Rule 1b is not sufficient. The problem is that, although Rule 1b was designed to solve a problem with ring-closure duplicate nodes,¹⁷ the rule as stated also applies to multiple-bond duplicate nodes. As such, we again have in Rule 1b the same issue involving multiple Kekulé structures as for Rule 1a. The problem involves cases such as shown in Figure 4, where the application of Rule 1b as




Figure 4. Two chemically equivalent Kekulé structures for the same molecular fragment.

currently recommended can lead to the inappropriate introduction of stereodescriptors. The two fragments are found to be identical by Rule 1a because the rings are carbocycles, so no atomic number averaging is needed. But in applying Rule 1b, we have a new issue, because the double bonds have different relationships to the phenolic OH.

Our discussion revolved around how to address this issue algorithmically: (a) Should we implement an averaging scheme as in Rule 1a for all six duplicate nodes? (b) Do we omit multiple-bond duplicate nodes from consideration—just assigning them “n/a” and skipping them entirely? (c) Should we assign the distance to the root of their sphere? (d) Do we assign the distance to the root of their attached atom? Our group decided that the first of these options was too complex, the second would lead to ambiguities in the algorithm, the third would not work, and the simplest and surest solution would be the last of these—to assign to a multiple-bond duplicate node the distance to the root of its corresponding attached atom, not its corresponding duplicated atom (Figure 5). Thus, this gives us the digraph for 2,2'-(hydroxymethylene)-diphenol (VS005), where numbers are distances to the root assigned for each node needed for Rule 1b (not the usual atomic numbers seen on digraphs in BB 2013). The specific representation of the bonding is no longer significant, and the center is found to have no descriptor, as expected.

There is an additional aspect of Rule 1b we suggest revising. The original statement of Rule 1b includes an additional criterion ranking any duplicate node higher than any node that is not a duplicate node for these purposes.¹⁷ This statement is omitted from BB 2013, ostensibly because it appears to be chemically irrelevant. However, it is important in the general statement of the rule topologically, and we suggest retaining it in CIP rules.

Thus, we propose the following new Rule 1b:

Rule 1b (proposed): Lower root distance precedes higher root distance, where “root distance” is defined: (a) in the case of ring-closure duplicate nodes as the sphere of the duplicated atom; (b) in the case of multiple-bond duplicate nodes as the sphere of the atom to which the duplicate node is attached; and (c) in all other cases as the sphere of the atom itself.

Note that the rule for assignment of root distance to ring-duplicate nodes is unchanged in this proposal. We note that an unresolved problem with Rules 1a and 1b involves the case of fully conjugated systems such as those shown in Figure 6 (VS215, VS217, and VS218). Each of these structures has at least two Kekulé representations, and yet they are not mancude-ring systems, so atomic number averaging is not formally called for by BB 2013 guidelines. Still, application of Rule 1a should not resolve the chirality. Rather, these descriptors must be decided by Rule 5, since the compound has two pseudoasymmetric centers. The solution to this issue, we propose, is to not focus on mancude systems at all. The essential feature is that there exists a fully conjugated cycle of double bonds of essentially any size.

Rule 2: Higher atomic mass number precedes lower.

The implementation problem in this case relates to comparisons where one atom has an isotope indicated and one does not, and also (again) when several alternative Kekulé structures are involved. The problem is that “mass number” is always an integer—the sum of the number of protons and neutrons in the nucleus. This leaves open the question as to what to do in the case where an atom with isotope number indicated is compared with an atom of the same element that




Figure 5. Analysis of VS005 using the current wording of Rule 1b (on the left) gives different root distances for the two branches, resulting in a higher priority for the ring on the left. The proposed solution (on the right) is to assign the root distances for multiple-bond duplicate nodes to be the root distance of their attached atom, not that of the atom they duplicate. Digraph numbers are distances to the root assigned for each node. *ns* indicates nonstereogenic.




Figure 6. Though not mancude-ring systems, these compounds have two or more equivalent Kekulé structures. They require atomic number averaging in Rule 1a and revised definition of root distance in Rule 1b in order to pass the decision to Rule 5, giving an (*s,s*) descriptor in each case.

has no isotope number indicated, thus referring to a natural composition for the element. Conflicting examples are given in BB 2013 section P-92.3 in relation to how to deal with this issue, where the term “mass number” is replaced with “atomic mass”.

The consensus of our group was to recommend changing the language of Rule 2 to be specific, using exact isotopic mass when an isotope is indicated, and atomic weight when it is not. In addition, we recommend explicitly using a mass of 0 for all duplicate nodes, since if there is a mass issue related to a node, it will always be found first in relation to its real atom before the duplicated atom is ever checked in Rule 2. We propose:

Rule 2 (proposed): Higher mass precedes lower mass, where mass is defined in the case of a duplicate node as 0, an atom with isotope indicated as its exact isotopic mass, and in all other cases as the element's atomic weight.

Atomic weights of the elements and isotopic abundances are IUPAC recommended values.^{18,19} Using this modification, the example in BB 2013 P-92.3 still holds: ⁸¹Br > Br > ⁷⁹Br. A second example in BB 2013 (p 1189 example 2), however, places ¹²⁵I > I, even though natural iodine is 100% I-127; we consider this an erratum in BB 2013; the corrected version should read ¹²⁵I < I = ¹²⁷I. In addition, C (12.011) > ¹²C (12.000), ¹⁶O (15.994) < O (15.999), and H(1.0079) > ¹H (1.0078).

The switch to atomic weight and exact isotope mass rather than integer mass number allows atoms of elements that are 100% one isotope naturally (⁹Be, ¹⁹F, ²³Na, ²⁷Al, ³¹P, ⁴⁵Sc, ⁵⁵Mn, ⁵⁹Co, ⁷⁵As, ⁸⁹Y, ⁹³Nb, ¹⁰³Rh, ¹²⁷I, ¹³³Cs, ¹⁴¹Pr, ¹⁵⁹Tb, ¹⁶⁵Ho, ¹⁶⁹Tm, ¹⁹⁷Au, ²⁰⁹Bi, ²³¹Pa, and ²³²Th) to be equivalent whether their isotope number is given explicitly or not, as is the case chemically. In addition, the elements Tc, Pm, Po, At, Rn,

Fr, Ra, Ac, and all elements with atomic number >92 have no natural abundance; their “atomic weight” found on the periodic table is just one of their integer isotope mass numbers, as though that isotope were 100% naturally abundant.

It may seem that the switch to exact isotope mass from integer isotope number might be difficult to implement, requiring access to a complete table of isotopes, but that is not the case. It turns out that we can use integer isotope mass numbers provided we take account of just four anomalies: ¹⁶O, ⁵²Cr, ⁹⁶Mo, and ¹⁷⁵Lu. These four isotopes are the only ones that have exact masses slightly below their element's atomic weight even though their mass number is above it. For example, the exact mass of ¹⁶O is 15.994, which is below the element's atomic weight of 15.999, even though its mass number (16) is higher. In practice, this is no problem. We simply use integer isotope numbers but reduce these four by 0.1 for the purpose of setting priorities. For example, since the atomic weight of oxygen is 15.999, when ¹⁷O is compared to O, there is no problem—17 > 15.999. But when ¹⁶O is specified, we use 15.9 instead of its actual value of 15.994. This allows ¹⁶O to have the required lower priority than “O” itself, with atomic weight 15.999. No other isotopes have this problem, and we can just use their unadjusted integer mass number as a surrogate for isotopic mass. In this way, there is never a need to check a table of exact isotope masses.

Once Rule 1b is passed, duplicate nodes should play no role in deciding configuration. The exclusion of duplicate nodes by assigning their mass to be zero guarantees this. It also removes an issue similar to the one discussed for Rules 1a and 1b, that different arrangements of double bonds in conjugates systems must not affect the configuration. For example, isotopically labeled alcohol VS007, shown in Figure 7, is achiral. To ensure this result, we simply assign all duplicate node masses to be zero, allowing the mass difference to be carried only by their corresponding duplicated atom.

Rule 3: When considering double bonds and planar tetraligand atoms, “seqcis” = “Z” precedes “seqtrans” = “E”, and this precedes nonstereogenic double bonds.

As mentioned above, generation of a complete digraph, including auxiliary descriptors, is required prior to Rule 3 (not just Rule 4a, as mentioned in BB 2013). This is because *seqcis* and *seqtrans* also describe double bonds that involve two diastereomorphic ligands on the same end (with stereo-descriptors RR and SR, for example). Inverting about a plane




Figure 7. VS007. Priority of an isotopically labeled atom in an aromatic ring in relation to Rule 1b must not depend upon the Kekulé structure used to depict the compound. To ensure this, we propose using zero for the mass of all duplicate nodes.

changes the comparison (to SS and RS, in this case), but this change does not reverse priorities.

Placement of Rule 3 before Rule 4a ensures that only enantiomorphic (*seqCis* and *seqTrans*) comparisons involving double-bonds and cumulenes with an odd number of double bonds are left to consider in Rules 4 and 5. From an implementation point of view, application of Rule 3 is simply the comparison of two ligand pairs, one pair on each end of an alkene or cumulene.

Rule 4a: *Chiral stereogenic units precede pseudoasymmetric stereogenic units, and these precede nonstereogenic units.*

That is, $(R \text{ or } S) > (r \text{ or } s)$, $(M \text{ or } P) > (m \text{ or } p)$, and $(seqCis \text{ or } seqTrans) > (seqcis \text{ or } seqtrans)$, and that all of these have higher priority than digraph nodes with no auxiliary descriptor. The purpose of Rule 4a is to ensure that all comparisons in Rule 4b and later are of the same general type: $R \text{ vs } S$, $M \text{ vs } P$, or $seqCis \text{ vs } seqTrans$ in Rules 4b; $r \text{ vs } s$ or $m \text{ vs } p$ in Rule 4c; $R \text{ vs } S$ or $M \text{ vs } P$ in Rule 5. In addition, application of Rule 4a guarantees that the lists of ranked descriptors that are being compared in Rule 4b are of equal length. Implementation of Rule 4a is straightforward and needs no further discussion.

Rule 4b: *When two ligands have different descriptor pairs, then the one with the first chosen like descriptor pair has priority over the one with a corresponding unlike descriptor pair.*

Rule 4b is by far the most difficult rule to comprehend and implement. One simplification is that although Rule 4b, as stated in BB 2013, refers to all possible mixes of R/S , M/P , and $seqCis/seqTrans$ descriptors, for implementation purposes, all auxiliary descriptors can be normalized by labeling them either R or S . For example, any of R , M , or $secCis$ can be assigned R for the purpose of processing Rules 4b. In this way, all discussion can be expressed in terms of “equal” or “not equal” to a reference R or S , rather than “like” vs “unlike”. When assigning $seqCis/seqTrans$ and M/P auxiliary descriptors, which involve multiple atoms, it is critical that an implementation assign those descriptors to the node that is closest to the root. Otherwise the second phase of Rule 4b may fail.

The process for ranking ligands in Rule 4b is a more complex process than for previous rules, involving a two-stage process. First, the nodes are reranked in a way that may cross digraph branches. Second, the nodes are scanned in rank order for auxiliary descriptor similarity to both R and S reference descriptors. The higher priority ligand is the one with the highest score after both of these comparisons are made. We have found the easiest way to conceptualize this algorithmically is to “read” each of the four rankings (two for each ligand) as a series of 0s and 1s, which can be implemented as an integer, an array, or a bit set. The examples in this paper use integers, though our different implementations actually use different representations. The basic idea is to create four lists for each pair of ligands that can be compared together. If there is no winner, we go on to the next rule.

Ranking of ligand nodes starts with the creation of a new priority criterion for each node. This criterion must incorporate the full path from the root to this node, including all Rule 4a priorities as well as the similarity or dissimilarity of the node’s descriptor to the reference. Pseudoasymmetric descriptors r , s , m , p , *seqcis*, and *seqtrans* are ignored in this process. As always, only previously identically ranked nodes are resolved.

The example in Figure 8 illustrates the process used in Rule 4b for VS262. Note that branches change order, depending

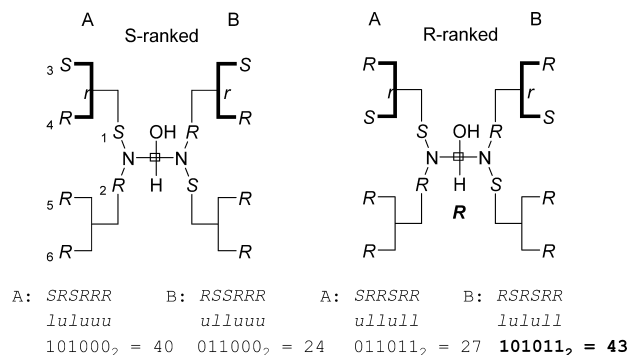


Figure 8. Numerical rankings of ligands A and B in Rule 4b for VS262 which requires using both S and R as the reference descriptor. Rule 4a has already sorted the main branches of the ligands, leaving only the bolded r branch for sorting. Ligand B, with the high score of 43, has the higher priority, and the descriptor is R .

upon the reference. In this case, some of the branch orderings have already been set, due to a Rule 4a comparison, $r \text{ vs } nonstereogenic$. The reading of the S -ranked ligand A, *SRSRRR* (read from centers 1–6, in order), is encoded as 101000 in base-2, giving a value of 40. Similar encoding gives 24 for S -ranked ligand B, 27 for R -ranked ligand A, and 43 for R -ranked ligand B. So ligand B, with a high score of 43, is given higher priority than ligand A, and the descriptor is R .

The rationale behind Rule 4b becomes clear upon inspection of the algorithm. First, we are giving preference to ligands that have the most similarities closest to the root—both RR and SS (each encoded as the number 3 by one of the reference options) will be selected in preference to either RS or SR (each of which will be encoded as 2 by one of the options and 1 by the other). And yet, identical readings or opposite readings overall—such as *RRSRRS* vs *SSRSSR*, which need to pass on to Rule 5—are not distinguished, as they will be encoded as the same two numbers with one or the other reference options.

An important facet of the ranking in Rule 4b is that identically ranked nodes can come from different branches of a ligand. So, for example, in Figure 9 we have a case (VS242), which has four “highest-ranked nodes” that must be sorted by R and S as a group. The result is that the two ligands cannot be distinguished in Rule 4b.

Note that all of the “rules for sorting” discussed in other works, such as, “The first step of this procedure is the critical choice of the first descriptor for each ligand,”²⁰ are unnecessary in terms of code implementation. These “critical choices”, such as determining the descriptor associated with the highest-ranking node or the descriptor that occurs the most in the set of highest-ranking nodes or sequentially evaluating both R and S as references, simply fall out of the mathematics of the numerical rankings described above. Thus, no such critical

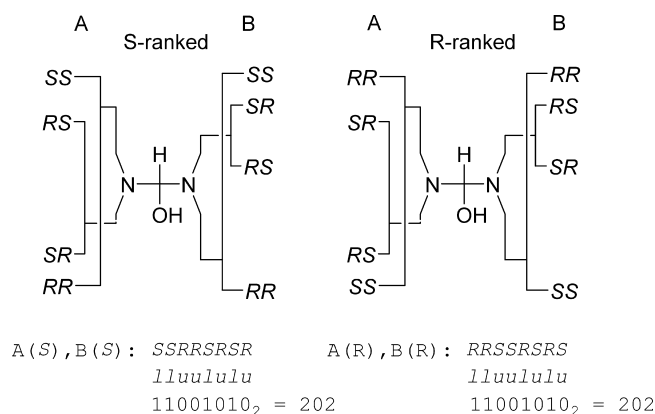


Figure 9. Ranking of VS242. In this case, ranking by references R and S involves crossing digraph lines, because in each case all four of the branches have “highest priority”. The ligands are not distinguished in Rule 4b.

choices need to be implemented, though, if they are, they might speed up the processing.

Rule 4c: “r” precedes “s” and “m” precedes “p”.

If a center passes Rule 4b undecided, it means that there are only three possibilities: (1) there is no ligand chirality; (2) two or more ligands have identical chirality descriptors; or (3) the two ligands each have sub-branches with opposite chirality. Rule 4c takes care of case 3, where we assign *r* over *s*, and *m* over *p*. The implementation of Rule 4c is a straightforward extension of the implementation of Rule 4a.

Rule 5: An atom or group with descriptor “R”, “M”, or “seqCis” has priority over its enantiomorph “S”, “P”, or “seqTrans”.

Rule 5 does a final check for enantiomorph ligands. Note that implementation of Rule 5 is not just a check of the lists generated using the procedure of Rule 4b, as priorities may have changed after application of Rule 4c. Consider the digraph in Figure 10, which shows the digraph of Figure 9, ranked by R- and S- reference for both Rule 4b and Rule 5. Here we see that after application of Rule 4c, the sorting is changed, and Ligand A has preference over Ligand B by Rule 5.

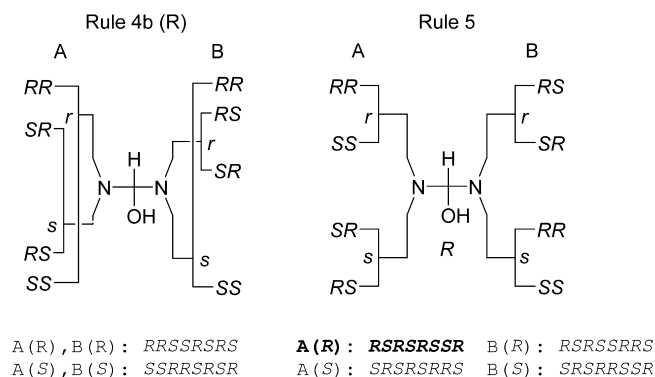


Figure 10. Same digraph as in Figure 9, for VS242, here also indicating the four pseudoasymmetric centers. In this case, sorting and listing of ligands gives different results after application of Rule 4c, which sorts both main ligand branches by *r* > *s*. The R-reference, Ligand A, with RSRSR... has higher priority than B, with RSRSS... Due to the fact that both asymmetric ligands are their own enantiomorph, the final descriptor is R rather than r.

If all ligands are finally distinguished after application of Rule 5, an additional test should be done to count the number of pairs of enantiomorph ligands. The final descriptor will be *r/s*, *m/p*, or, in the case of alkenes, *seqCis/seqTrans*, if and only if this number is one (Figure 11); otherwise, it will be *R/S*, *M/P*, or *seqcis/seqtrans* (*Z/E*).




Figure 11. (a) VS236, a typical center decided by Rule 5, with one pair of enantiomorph ligands; (b) VS244, a doubly enantiomorph center decided by Rule 5 to be S; (c) VS229, an alkene with one pair of enantiomorph ligands given the descriptor *seqTrans* (indicated as “e” in the validation suite) not *seqtrans (E)* because, upon reflection through a plane, it changes descriptor; and (d) VS245, an alkene with two pairs of enantiomorph ligands given the usual *seqtrans* (or *E*) descriptor. Similar tests need to be made for both odd- and even-cumulenes.

In terms of implementation, the criterion for pseudoasymmetry at a tetrahedral center is that, when comparing otherwise identical ligands, there is an odd number of pairs that reverses priority when comparing like/unlike sequences using an S reference vs using an R reference. So, for example, in Figure 11a, the right-hand R ligand has higher priority with the R-reference, but the left-hand S ligand has higher priority with the S-reference. The priority switches, and we have the normal outcome for Rule 5—pseudoasymmetric *r*. But in Figure 11b, priority switches twice, so the result is S, not s.

In Figure 10, we have a different story. Sorting by R gives A > B. But sorting by S also gives A > B—no switch! A naive application of Rule 5, only checking the two R-reference like/unlike lists, would have assigned *r* to that center. However, this center is asymmetric, not pseudoasymmetric, because each ligand is its own enantiomorph. The ultimate descriptor will be R, not r. The analysis is the opposite for alkenes and even-atom cumulenes (Figure 11c and d). Thus, if the pair on only one end of the alkene reverses priority when using the S reference vs using the R reference (Figure 11c), then the result is the asymmetric *seqCis* or *seqTrans*; if neither or both pairs reverse priority (Figure 11d), then the result is the pseudoasymmetric *seqcis* or *seqtrans*.

Alternatively, one can simply repeat the entire process of assigning descriptors using the structure that is reflected through a mirror plane. If the resulting descriptor is reversed, then the final descriptor is *R/S*, *M/P*, or *seqCis/secTrans*; if not, then *r/s*, *m/p*, or *seqcis/seqtrans*.

Proposal for Rule 6: Spiro and Other Axially Symmetric Compounds. Early on in the development of the CIP system,³ it was recognized that certain cases involving C₂, D₂, and C₃ point groups require additional consideration for assignment of stereodescriptors; an S₄ case was added later³ (Figure 12). Specifically, an algorithm must distinguish between both enantiomers of compounds VS285, VS281, and VS283, and yet still deliver no stereodescriptors for cubane (VS009) and the S₄ compound VS012. Only simple spiro structure VS285 is mentioned in BB 2013, in section P-93.5.3.

A relatively simple additional ninth Sequence Rule—“Rule 6”—takes care of all such cases:




Figure 12. Five compounds for which the stereodescriptor is decided by Rule 6, along with their point group and final descriptor.

Rule 6 (proposed): An undifferentiated reference node has priority over any other undifferentiated node.

Cases to be considered by Rule 6 are identified by having, after application of Rule 5, two pairs of identical ligands or three or four identical ligands. Thus, in the first case in Figure 12, we have two identical amino ligands and two identical keto ligands; in the second and fifth cases, we have four identical ligands. In the third and fourth cases, we have three identical ligands.

The solution for all such cases (first proposed in 1966² for specific examples) is simply to select one node of any one of the undistinguished ligands for promotion to higher rank (Figure 13). Basically, by arbitrarily breaking the symmetry in




Figure 13. Two examples of the analysis of compounds by Rule 6, which sets Node 1 to be higher priority than Node 2. This single change decides also the priority $3 > 4$, due to the presence of a ring connection from Node 3 back to Node 1 and from Node 4 back to Node 2.

this way, the problem is immediately resolved upon inspection of the digraph. We note for double spiran VS281, the result of application of the proposed Rule 6 generates the opposite descriptor to the one assigned previously.^{2,3} We believe this is a due to either a misassignment or a typographical error in that publication.

Two outcomes of Rule 6 are possible:

- After application of Rule 6, there are still two undistinguished ligands. Such will be the case, for example, with simple acyclic compounds, such as CH_2Cl_2 or CHCl_3 . The center remains without descriptor.
- After application of Rule 6, all ligands are distinguished. The center receives a descriptor. Such will be the case only for compounds that have rings that involve the root atom and three or more ligands. A full application of Rule 6 tests all possible promotions for the same results, though this is necessary only for certain symmetries. Any matching *R* and *S* pairs are ignored; if a descriptor remains, it is valid.

In terms of implementation of Rule 6, it may be noted that a descriptor will only be assigned by Rule 6 when there are three or more paths through the structure leading back to the root atom. Thus, a simple test that indicates fewer than three

duplicate nodes with root distance 0 is sufficient for skipping Rule 6.

Note that Rule 6 must be applied in the determination of all auxiliary descriptors as well as in the final root-node determination, as for all other Sequence Rules, because it is possible for an auxiliary descriptor to result from C_3 symmetry (e.g., VS300). Even if it is determined that Rules 4a–5 can be skipped due to the lack of auxiliary descriptors, Rule 6 should not be skipped without further evidence that outcome a is the only possibility.

Rule 6 also allows assignment of stereodescriptors for centers with axial symmetry, such as those involving allenes and biphenyls. No further special consideration of high-symmetry compounds or groups is necessary. Interestingly, an easy extension of Rule 6 solves the heretofore unresolved issue of “in/out” configuration such as shown in Figure 14. These




Figure 14. Stereodescriptors for (*r,r*)- and (*s,s*)-bicyclo[2.2.2]octane which arise from allowing a recalculation of auxiliary descriptors in the processing of Rule 6. Similar descriptors are found for all-out and in/out tetrahedrane, cubane, adamantane, and related structures.

compounds can be treated successfully simply by taking into account the temporary reassignment of auxiliary descriptors after each promotion of a node. However, we stop short of recommending this modification at this time, as it only applies to rather esoteric structures mostly of only theoretical interest. Its implementation would result in additional (unnecessary) descriptors for common all-*r* “all-out” structures, many of which are known compounds and have already been arbitrarily assumed to have no stereodescriptors.

Implementations. The CIP rules defined in BB 2013 and all relevant publications were analyzed in the development of algorithmic procedures for generation of stereodescriptors. Following this analysis, it was concluded that several refinements are necessary for the CIP rules to ensure generation of the descriptors described in BB 2013.

Specific software implementations are listed in Table 1, with detailed implementation characteristics given in Table 2. The four implementations were developed independently. Two are open-source; all are freely available as executable versions. As can be seen, all implementations have limitations. However, outside of a few specific limitations, all four implementations give the same results for all compounds in the validation set and agree in all cases with descriptors provided in BB 2013 or its errata.

Table 1. Software Availability, Features, and Limitations

	Jmol (R.M.H., http://jmol.sourceforge.net , https://chemapps.stolaf.edu/jmol/jsmol/cip.htm)	
Availability	Open-source; both Java and JavaScript versions (v. 14.29.15)	
Features	Visual; interactive 3D models; easy building and adaptation of models (click atom to invert chirality); can be made part of any web page in the form of a JSmol app, run on a server using JmolData.jar, or run as a stand-alone Java application as Jmol.jar. The Jmol implementation of CIP rules leverages Jmol's substantial features in the area of 3D molecular data handling and visualization, including Jmol conformational SMARTS ²¹ searching ²² to effectively identify potential atropisomer and helicene chirality. Scripting provides easy annotation of 3D interactive models with descriptors (<i>LABEL %[chirality]</i>) and the CIP Rules leading to those (<i>LABEL %[cipRule]</i>), atom selection (<i>select chirality == "R"</i>), and structured return of CIP information (<i>calculate chirality; print_M.CIPInfo</i>). 2D-to-3D and SMILES-to-3D conversion are provided by the NCI/CADD Chemical Resolver. ²³	
	Balloon (M.J.V., http://users.abo.fi/mivainio/balloon)	
Availability	Closed-source, freeware; C++ (v1.6.7)	
Features	Balloon is a command-line program for 2D to 3D conversion using distance geometry and conformational sampling using a genetic algorithm. The chirality perception code, amounting to around 2K lines, is implemented in C++. The algorithm searches the input molecular graph for possible stereogenic centers and generates a digraph for each center. Conformational sampling is possible with input chirality retained.	
	ACD/Name and ACD/ChemSketch (A.Y., http://www.acdlabs.com/products/draw_nom)	
Availability	Closed-source; freeware and commercial versions; Delphi (v. 14.05)	
Features	ACD/Name generates chemical names for structures indicating stereodescriptors in names in accord with IUPAC recommendations. It specifies stereoconfiguration with <i>E/Z</i> , <i>R/S</i> , and <i>P/M</i> descriptors supporting tetrahedral and trigonal pyramidal centers, cumulene and atropisomer axial chirality. It supports both 2D structures with stereobonds and 3D structures. ACD/ChemSketch is a desktop full featured chemical drawing program allowing stereodescriptor generation using ACD/Name's stereochemistry procedures.	
	Centres (J.W.M., http://www.github.com/simolecule/centres)	
Availability	Open-source; 2-Clause BSD, Java (v.1.0)	
Features	Centres is a library specifically for CIP labeling. It provides an abstract API allowing integration on top of multiple back-end toolkits (CDK, ^{24–27} OPSIN, ²⁸ JChem Base ²⁹). The underlying toolkit handles the file processing and describes the stereochemistry as a data structure. Centres then generates CIP descriptors.	

Table 2. Features and Limitations of the Implementations

	features	Jmol	Balloon	ACD ChemSketch	Centres	examples
software details	language	Java/JS	C/C++	Delphi	Java	
	free to use	Y	Y	Y ^a	Y	
	open-source	Y	N	N	Y	
priority rules	Rule 1a	Y	Y	Y	Y	VS013-170
	Rule 1b ^b	Y	Y	Y	Y	VS171-174
	Rule 2 ^b	Y	Y	Y	Y	VS175-187
	Rule 3	Y	Y	Y	Y	VS188-195,246-248
	Rule 4a	Y	Y	Y	Y	VS249-251,269-272,277,278
	Rule 4b	Y	Y	Y	Y	VS196-204,252-263,265,268-272,279
	Rule 4c	Y	Y	Y	Y	VS273-279,296-298
	Rule 5	Y	Y	Y	Y	VS205-300
	Rule 6 ^b	Y	Y	Y	Y	VS280-300
features	no stereo	Y	Y	Y	Y	VS001-009,012
	<i>seqcis/seqCis</i> distinction	Y ^c	Y ^c	Y ^c	Y	VS229,246-248,299
	large cyclic π systems	Y	Y ^d	Y	Y	VS215-218
geometry support	odd-cumulene (chirality axis)	Y	Y	Y	Y	VS078,079,120,141,144,166,231,232,243,287
	even-cumulene (planar)	Y	Y	Y	Y	VS063.118.135,154,164
	atropisomer (chirality axis)	Y	N	Y	Y	VS023,055,057,072,073,086,158
	helical chirality	Y	N	N	N	VS010,011
	fullerene chirality	N	N	N	N	
	chirality plane	N	N	N	N	
	inorganic configuration index	N	Y	N	Y	

^aStructure size restriction in freeware. ^bProposed Sequence Rules revision or addition. ^cReports *z/e* for *seqCis/seqTrans* and *Z/E* for *seqcis/seqtrans*.

^dUses aromaticity flags in lieu of full detection.

METHODS—VALIDATION

Key to our development process was the production of a robust validation suite that could be used by any developer to ensure that a CIP implementation follows the rules to whatever degree that software claims to implement them. Though scattered attempts have been made to develop such model collections, primarily for in-house testing of various software packages, the Supporting Information for this paper includes the first fully tested openly available collection of models that can test the full range of issues presented in the CIP Sequence Rules. The 300 models are in annotated structure data file (SDF) format, providing both 2D and 3D models, as well as

SMILES.²² Standard SDF data annotations provide reported (or corrected) descriptors for all relevant models in Chapter 9 of BB 2013, along with a significant number of additional “challenges” for developers. Accompanying this collection is a table that correlates the specific test models with specific sections of BB 2013 Chapter 9 as well as specific aspects of the CIP rules that are targeted by this particular test.

The four implementations generate the same stereodescriptors for each model in the validation set. The analysis largely agrees with the descriptors provided in BB 2013, differing only in relation to a few already-identified minor

analysis or typographical errors in the published version of BB 2013.

CONCLUSIONS

We have developed and implemented specific algorithms that handle all nuances of all nine of the Sequence Rules for tetrahedral centers and double bond configuration, two of which are open-source Java applications (Centres and Jmol, which also has a JavaScript equivalent). All are freely available in binary format. We have developed a robust testing suite for algorithms that implement the CIP Sequence Rules. This suite covers a wide variety of possible issues that might arise in the course of any algorithm development.

In addition, in the process of this work, we have identified and addressed a number of issues with the eight Sequence Rules as presented in *Nomenclature of Organic Chemistry: IUPAC Recommendations and Preferred Names 2013* (1a, 1b, 2, 3, 4a, 4b, 4c, and 5) and propose simple solutions, including a new Rule 6. In particular, we provide a definitive method of implementing Kekulé averaging in Rules 1a, 1b, and 2. We propose a simple solution to working with isotope masses in Rule 2. We argue for the necessity for fully elaborated auxiliary descriptors prior to Rule 3 and describe the order in which they must be generated. We demonstrate a simple way to implement the *like/unlike* analysis in Rules 4b and 5. Most significantly, we propose three modifications to the Sequence Rules:

Proposed New Rule 1b: Lower root distance precedes higher root distance, where “root distance” is defined: (a) in the case of ring-closure duplicate nodes as the sphere of the duplicated atom; (b) in the case of multiple-bond duplicate nodes as the sphere of the atom to which the duplicate node is attached; and (c) in all other cases as the sphere of the atom itself.

Proposed New Rule 2: Higher mass precedes lower mass, where mass is defined in the case of duplicate nodes as 0, atoms with isotope indicated as their exact isotopic mass, and in all other cases, as their element’s atomic weight.

Proposed New Rule 6: An undifferentiated reference node has priority over any other undifferentiated node.

ASSOCIATED CONTENT

Supporting Information

The Supporting Information is available free of charge on the ACS Publications website at DOI: 10.1021/acs.jcim.8b00324.

Collection of models that can test the full range of issues presented in the CIP Sequence Rules (ZIP)

Spreadsheet correlating VS-number, SMILES, stereochemical descriptors, BB 2013 section, type of stereochemical unit(s), relevant CIP Rule(s), and Figure numbers (XLS)

AUTHOR INFORMATION

Corresponding Author

*E-mail: hansonr@stolaf.edu.

ORCID

Robert M. Hanson: 0000-0001-5411-2356

Sophia Musacchio: 0000-0001-8191-2869

John W. Mayfield: 0000-0001-7730-2646

Mikko J. Vainio: 0000-0001-9005-6590

Andrey Yerin: 0000-0003-1523-0020

Dmitry Redkin: 0000-0003-3420-6931

Notes

The authors declare no competing financial interest.

ACKNOWLEDGMENTS

The authors thank the International Union of Pure and Applied Chemistry (IUPAC) for permission to use and provide chemical structures from BB 2013 as part of the validation suite. R.M.H. acknowledges Trâm Bui for her early work with the structure database. A.Y. thanks the European Commission Taxation and Customs Union (TAXUD/2007/CC/089 and TAXUD/2012/CC/119) for funding ACD/Laboratories to digitize chemical structures and names from BB 2013 into SDF format, some of which were used in the validation suite.

REFERENCES

- (1) Cahn, R. S.; Ingold, C. K.; Prelog, V. The specification of asymmetric configuration in organic chemistry. *Experientia* **1956**, *12*, 81–94.
- (2) Cahn, R. S.; Ingold, C.; Prelog, V. Specification of Molecular Chirality. *Angew. Chem., Int. Ed. Engl.* **1966**, *5*, 385–415.
- (3) Prelog, V.; Helmchen, G. Basic Principles of the CIP-System and Proposals for a Revision. *Angew. Chem., Int. Ed. Engl.* **1982**, *21*, 567–583.
- (4) Favre, H. A.; Powell, W. H. Chapter P-9. Specification of Configuration and Conformation. In *Nomenclature of Organic Chemistry: IUPAC Recommendations and Preferred Names 2013*; Royal Society of Chemistry: Cambridge, UK, 2013; pp 1156–1292.
- (5) Mata, P.; Lobo, A. M.; Marshall, C.; Johnson, A. P. The CIP Sequence Rules: Analysis and proposal for a revision. *Tetrahedron: Asymmetry* **1993**, *4*, 657–668.
- (6) Mata, P.; Lobo, A. M.; Marshall, C.; Johnson, A. P. Implementation of the Cahn-Ingold-Prelog System for Stereochemical Perception in the LHASA Program. *J. Chem. Inf. Model.* **1994**, *34*, 491–504.
- (7) Mayfield, J.; Lowe, D.; Sayle, R. CINF 17: Comparing Cahn-Ingold-Prelog Rule Implementations: The need for an open CIP. <https://www.slideshare.net/NextMoveSoftware/cinf-17-comparing-cahningoldprelog-rule-implementations-the-need-for-an-open-cip> (accessed April 2, 2018).
- (8) Corrections, Revisions and Extensions for the Nomenclature of Organic Chemistry—IUPAC Recommendations and Preferred Names 2013 (The IUPAC Blue Book). https://iupac.org/projects/project-details/?project_nr=2015-052-1-800 (accessed April 2, 2018).
- (9) BlueObelisk. <https://blueobelisk.github.io> (accessed April 2, 2018).
- (10) Jmol. <http://jmol.sourceforge.net> (accessed April 2, 2018).
- (11) May, J. W. *Cheminformatics for genome-scale metabolic reconstructions (doctoral thesis)* **2015**, DOI: 10.17863/CAM.15987.
- (12) ACD/ChemSketch Freeware. http://www.acdlabs.com/products/draw_nom/draw/chemsketch/ (accessed 2018).
- (13) Vainio, M. J.; Johnson, M. S. Generating conformer ensembles using a multiobjective genetic algorithm. *J. Chem. Inf. Model.* **2007**, *47*, 2462–2474.
- (14) Puranen, J. S.; Vainio, M. J.; Johnson, M. S. Accurate conformation-dependent molecular electrostatic potentials for high-throughput in silico drug discovery. *J. Comput. Chem.* **2009**, *31*, 1722–1732.
- (15) Ziegler, F., Study Aids Chemistry220js. <http://ursula.chem.yale.edu/~chem220/chem220js/StudyAids.html#Stereochemistry> (accessed March 31, 2018).
- (16) Mancude-ring systems. <https://goldbook.iupac.org/html/M/M03695.html> (accessed March 31, 2018).
- (17) Custer, R. H. Mathematical statements about the revised CIP-system. *Match* **1986**, *21*, 3–31.

(18) Meija, J.; Coplen, T. B.; Berglund, M.; Brand, W. A.; De Bièvre, P.; Gröning, M.; Holden, N. E.; Irrgeher, J.; Loss, R. D.; Walczyk, T.; Prohaska, T. Atomic weights of the elements 2013 (IUPAC Technical Report). *Pure Appl. Chem.* **2016**, *88*, 265–291.

(19) Rosman, K. J. R.; Taylor, P. D. P. Isotopic compositions of the elements 1997 (Technical Report). *Pure Appl. Chem.* **1998**, *70*, 217–235.

(20) Mata, P.; Lobo, A. M. The Cahn, Ingold and Prelog System: eliminating ambiguity in the comparison of diastereomorphic and enantiomorphic ligands. *Tetrahedron: Asymmetry* **2005**, *16*, 2215–2223.

(21) Daylight Theory Manual. <http://www.daylight.com/dayhtml/doc/theory/index.html> (accessed March 31, 2018).

(22) Hanson, R. M. Jmol SMILES and Jmol SMARTS: Specifications and Applications. *J. Cheminf.* **2016**, *8*, 50.

(23) Sitzmann, M.; Ihlenfeldt, W. D.; Nicklaus, M. C. NCI/CADD: Open-access chemical structure web platform. In *Abstr. Pap. Am. Chem. Soc.*, 2010; Vol. 239.

(24) Willighagen, E. L.; Mayfield, J. W.; Alvarsson, J.; Berg, A.; Carlsson, L.; Jeliaskova, N.; Kuhn, S.; Pluskal, T.; Rojas-Chertó, M.; Spjuth, O.; Torrance, G.; Evelo, C. T.; Guha, R.; Steinbeck, C. The Chemistry Development Kit (CDK) v2.0: atom typing, depiction, molecular formulas, and substructure searching. *J. Cheminf.* **2017**, *9*, 33.

(25) Steinbeck, C.; Hoppe, C.; Kuhn, S.; Floris, M.; Guha, R.; Willighagen, E. Recent Developments of the Chemistry Development Kit (CDK) - An Open-Source Java Library for Chemo- and Bioinformatics. *Curr. Pharm. Des.* **2006**, *12*, 2111–2120.

(26) May, J. W.; Steinbeck, C. Efficient ring perception for the Chemistry Development Kit. *J. Cheminf.* **2014**, *6*, 3.

(27) Steinbeck, C.; Han, Y.; Kuhn, S.; Horlacher, O.; Luttmann, E.; Willighagen, E. The Chemistry Development Kit (CDK): An Open-Source Java Library for Chemo- and Bioinformatics. *J. Chem. Inf. Comput. Sci.* **2003**, *43*, 493–500.

(28) Lowe, D. M.; Corbett, P. T.; Murray-Rust, P.; Glen, R. C. Chemical Name to Structure: OPSIN, an Open Source Solution. *J. Chem. Inf. Model.* **2011**, *51*, 739–753.

(29) ChemAxon. JChem Base. <https://chemaxon.com/products/jchem-engines> (accessed May 9, 2018).