Bioorganic & Medicinal Chemistry 20 (2012) 5343–5351

Contents lists available at SciVerse ScienceDirect

Bioorganic & Medicinal Chemistry

journal homepage: www.elsevier.com/locate/bmc

Softening the Rule of Five—where to draw the line? ☆

Joachim Petit [a], Nathalie Meurice [a], Christine Kaiser [a,b], Gerald Maggiora [a,b,*]

[a] Translational Genomics Research Institute, 445 N 5th Street, Suite 600, Phoenix, AZ 85004, United States
[b] Department of Pharmacology and Toxicology, University of Arizona College of Pharmacy, Tucson, AZ 85721, United States

ARTICLE INFO

Article history:
Available online 17 December 2011

Keywords:
Lipinski Rules
Soft Rule of Five
Utility functions
Analytic hierarchy process
Drug likeness
Decision theory
Molecular properties

ABSTRACT

In order to improve the discovery and development of new drugs, a broad effort is being made to assess the 'drug-like' properties of molecules in early stages of the discovery-research process. Although there are numerous approaches to this problem, perhaps the simplest and most widespread one is that developed by Chris Lipinski and his co-workers at Pfizer, which is generally referred either as the Lipinski Rules or the Rule of Five (ROF). The ROF is based on four properties of molecules, namely, molecular weight (MW), log*P*, number of hydrogen bond donors (HBD), and the number of hydrogen bond acceptors (HBA). A 'flag' is set if the value of a given property exceeds the chosen threshold value for that property—MW 500 Da, log*P* 5, the number of HBDs 5, and the number of HBAs 10. Each flag corresponds to an ROF violation. The total number of violations is the ROF-Score, which lies between '0' and '4'. Molecules with ROF-Scores greater than one are considered to be marginal for further development. The difficulty with this approach is that two molecules with nearly identical property values can, nonetheless, possess ROF-Scores that can differ by two or more. Thus, one molecule could be considered for further studies while the other, nearly identical molecule (in terms of its four ROF properties), would most likely not be. This problem arises because of the sharp thresholds imposed by the present formulation of the ROF, which is based upon classical sets. In the current work an alternative approach based on the use of utility functions, within the framework of the analytic hierarchy process (AHP), are employed to 'soften' the sharp boundaries inherent in classical sets. This provides a more realistic assessment of compounds in terms of their potential suitability in drug-discovery research programs.

1. Introduction

In order to improve the discovery and development of new drugs, a broad effort is being made to assess the 'drug-like' properties of molecules in early stages of the discovery-research process. Although there are numerous approaches to this problem,[1–6] perhaps the simplest and most widespread one is that developed by Chris Lipinski and his colleagues at Pfizer,[7] which is generally referred either as the Lipinski Rules or the Rule of Five (ROF). The ROF is based on four properties of molecules, namely, molecular weight (MW), log*P*, number of hydrogen-bond donors (HBD) taken as equivalent to the number of –OH and –NH groups, and the number of hydrogen-bond acceptors (HBA) taken as equivalent to the number of oxygen and nitrogen atoms. A 'flag' is set if a molecule's MW is greater than 500, its log*P* is greater than 5, the number of its HBDs exceeds 5 and the number of its HBAs exceeds 10. Because the values of the decision points for all of the property values are multiples of five, the above set of rules has been called the 'Rule of Five.'

The total number of violations is the ROF-Score, which lies between '0' and '4'. Molecules with ROF-Scores greater than one are considered to be marginal for further development, although as pointed out by Lipinski and co-workers.[7] such molecules should not necessarily be removed from further consideration. Rather, they should be deprioritized in the discovery research process. Lastly, it is well known that many drugs violate the ROF, but this is not a serious issue since it was not originally designed as a tool for assessing drug likeness. Nevertheless, its common usage for this purpose has, de facto, made it so in practice.

Although the ROF has proved quite useful, it has some glaring weaknesses; addressing them is the subject of this paper. Basically, the two major weaknesses are the equal weight given to each of the rules and the sharp boundary that marks the violation of a given rule. For example, consider two molecules, one with molecular weight 501 Da and one with a molecular weight of 499 Da. Both molecules have very nearly the same molecular weights, but the former violates the molecular weight rule while the latter does not. This clearly indicates the need for an approach that is free

☆ This manuscript contains some material presented in the 2008 Herman Skolnik Award address by G.M. Maggiora at the American Chemical Society Meeting in Philadelphia, PA, August 18, 2008.

* Corresponding author.
E-mail address: gerry.maggiora@gmail.com (G. Maggiora).

0968-0896/$ - see front matter © 2011 Elsevier Ltd. All rights reserved.
doi:10.1016/j.bmc.2011.11.064

from the difficulties caused by the sharp boundaries that characterize violations of the rules in the current approach. This paper will address both of the weaknesses of the ROF using a procedure based on the analytic hierarchy process (AHP).[8] AHP is a decision theoretic method with widespread applications in many fields,[9] although few applications in biomedical science research (cf. Ref. 10). The goal of the paper is not a definitive study of the specific form that the methodologies described here will ultimately take. Rather, it is to introduce scientists to the softer approach proposed in this work. A detailed analysis of the optimum form of these methodologies will be dealt with in a future study.

The AHP provides a means for ranking a set of criteria based upon pairwise comparisons among them. Although the AHP is a general theory for decision making based upon a hierarchical decision structure, only a single, non-hierarchical decision structure is required in the present work since only a single set of criteria, namely the four ROF rules, is required (vide supra). Application of the AHP to this case will provide a set of weights for the four rules that will indicate the relative importance of each of the rules. An approach similar in form to the present work has been developed by Wager et al. to treat properties of CNS drugs.[11]

The paper will first provide a brief explanation of the AHP, the key methodology employed in this work. Following that several simple examples will be presented illustrating how softening the sharp decision rules of the classical ROF can provide a more realistic assessment of a set of molecules. This will be followed by an application of the soft ROF method to a sample of the DrugBank Library, a publicly available compound database of more than 4000 drugs and related compounds. The paper will conclude with a summary of the method and some conclusions that may be drawn from the examples considered in this work.

2. Methodology

2.1. The analytic hierarchy process

The AHP is the theoretical foundation of the method applied in this work. Thomas Saaty developed it in the 60's as an intuitive and fundamental approach to decision theory.[8] It is designed to cope with both objective and subjective data, is based on *pairwise comparisons*, allows for inconsistency of judgments, and provides a means for improving the consistency of judgments. The present work is aimed at applying the AHP to develop a softer and more robust version of the classical ROF, which *classifies* molecules into five discrete categories each of which depends on the number of times (0, 1, 2, 3, or 4) that the property values of a given molecule exceed their associated ROF thresholds. Each time a property value exceeds its ROF threshold a flag is set. If a molecule has two or more flags set (i.e., ROF violations) it is considered marginal for further drug development.

In contrast, the new method generates a *linear order* of molecules based on a continuous score that accounts for the combined degree that each of the ROF properties exceeds its threshold value. The constituents of the AHP decision system include the set of four ROF properties.

$$P = \{P_{\mathrm{MW}}, P_{\mathrm{LogP}}, P_{\mathrm{HBD}}, P_{\mathrm{HBA}}\} \tag{2.1}$$

and the set of n molecules

$$M = \{M_1, M_2, \ldots, M_n\} \tag{2.2}$$

each of which is evaluated with respect to the properties in P.

The relative importance of the properties is determined through pairwise comparisons (vide infra), which are based on an assessment of the 'importance' of one property to another. *As used here, importance is a general term that represents the importance, preference, desirability, dominance, likelihood, or whatever comparable term is appropriate*. Comparative assessments between properties can be subjective as well as objective. For example, when considering molecules for pre-clinical development how important are their solubilities compared to their molecular weights? While there is data to support various points of view on this question, a definitive answer is not currently available. Hence, the relative importance of the two properties is somewhat subjective in nature, although not arbitrary since it relies on the expert knowledge of those carrying out the study.

A key element of the AHP is the 4×4-*dimensional comparison matrix* $\mathbf{C}$, which characterizes the set of pairwise comparisons among the four ROF properties:

$$\mathbf{C} = \begin{bmatrix} C_{\mathrm{MW,MW}} & C_{\mathrm{MW,LogP}} & C_{\mathrm{MW,HBD}} & C_{\mathrm{MW,HBA}} \\ C_{\mathrm{LogP,MW}} & C_{\mathrm{LogP,LogP}} & C_{\mathrm{LogP,HBD}} & C_{\mathrm{LogP,HBA}} \\ C_{\mathrm{HBD,MW}} & C_{\mathrm{HBD,LogP}} & C_{\mathrm{HBD,HBD}} & C_{\mathrm{HBD,HBA}} \\ C_{\mathrm{HBA,MW}} & C_{\mathrm{HBA,LogP}} & C_{\mathrm{HBA,HBD}} & C_{\mathrm{HBA,HBA}} \end{bmatrix}. \tag{2.3}$$

Each element of $\mathbf{C}$, which is always taken to be greater than zero, can be interpreted as the relative importance of the π-th property over the π'-th property (vide supra), where $\pi,\pi' = \{\mathrm{MW}, \mathrm{Log}P, \mathrm{HBD}, \mathrm{HBA}\}$. A significant feature of the elements of the $\mathbf{C}$ matrix is their reciprocal character, namely,

$$C_{\pi,\pi'} = \frac{1}{C_{\pi',\pi}} \tag{2.4}$$

and thus $C_{\pi,\pi} = 1$. As an example, suppose the π-th property is twice as important as the π'-th property, then the π'-th property will be half as important as the π-th property.

Consistency of judgments is related to the notion of transitivity in set theory. Suppose the π-th property is three times as important as the π'-th property, and suppose the π'-th property is twice as important as the π''-th property, then for consistency the π-th property should be six times as important as the π''-th property. Hence, for consistency to obtain the product of the elements of the comparison matrix must satisfy $C_{\pi,\pi'} \cdot C_{\pi',\pi''} = C_{\pi,\pi''}$. However, as has been pointed out by Saaty, consistency need not be satisfied rigorously for an AHP analysis to be useful.[8]

Because of a desire to maintain consistency, an eigenvalue formulation was developed by Saaty,

$$\mathbf{C}\ \mathbf{v}^{\max} = \lambda_{\max}\mathbf{v}^{\max} \tag{2.5}$$

where $\lambda_{\max}$ is the maximum eigenvalue of $\mathbf{C}$ and $\mathbf{v}^{\max}$ is its corresponding eigenvector. In matrix form Eq. 2.5 is given by

$$\begin{bmatrix} 1 & C_{\mathrm{MW,LogP}} & C_{\mathrm{MW,HBD}} & C_{\mathrm{MW,HBA}} \\ (C_{\mathrm{MW,LogP}})^{-1} & 1 & C_{\mathrm{LogP,HBD}} & C_{\mathrm{LogP,HBA}} \\ (C_{\mathrm{MW,HBD}})^{-1} & (C_{\mathrm{LogP,HBD}})^{-1} & 1 & C_{\mathrm{HBD,HBA}} \\ (C_{\mathrm{MW,HBA}})^{-1} & (C_{\mathrm{LogP,HBA}})^{-1} & (C_{\mathrm{HBD,HBA}})^{-1} & 1 \end{bmatrix} \tag{2.6}$$

$$\cdot \begin{bmatrix} v_{\mathrm{MW}}^{\max} \\ v_{\mathrm{LogP}}^{\max} \\ v_{\mathrm{HBD}}^{\max} \\ v_{\mathrm{HBA}}^{\max} \end{bmatrix} = \lambda_{\max} \cdot \begin{bmatrix} v_{\mathrm{MW}}^{\max} \\ v_{\mathrm{LogP}}^{\max} \\ v_{\mathrm{HBD}}^{\max} \\ v_{\mathrm{HBA}}^{\max} \end{bmatrix}$$

It can be shown using the Perron-Frobenius Theorem[12] that since $\mathbf{C}$ is a 4×4-dimensional positive, reciprocal matrix, $\lambda_{\max}$ and the components of the corresponding eigenvector are all positive. The importance or 'weight' of the π-th component of $\mathbf{v}^{\max}$, which corresponds to the π-th property, can be obtained by normalizing its value by the sum of the values of all of the components.

$$w_\pi = \frac{v_\pi^{\max}}{\sum\limits_{\pi'} v_{\pi'}^{\max}}, \quad \sum_{\pi} w_\pi = 1. \tag{2.7}$$

Ordering the components in accordance with the values of their weights yields a linear ordering of the properties with respect to their relative importances. As shown by Saaty,[8] weights obtained in this manner lead to a consistent comparison matrix C whose maximum eigenvalue satisfies $\lambda_{max} \geqslant 4$. Equality obtains only when the **C** matrix is completely consistent. This relationship is the basis for the development of a *consistency index*.

$$\mathrm{CI} = \frac{\lambda_{max} - 4}{3}, \tag{2.8}$$

where the term in the denominator corresponds to the degrees of freedom in the set of weights. Thus, the larger the CI, the greater the inconsistency of the **C** matrix. Saaty normalized the CIs by combining them with what he called the *random consistency index* RI, which is obtained by sampling a very large number of randomly generated 4×4-dimensional **C** matrices and computing the mean of their CIs. A *consistency ratio* CR, which provides a measure of the degree of consistency of the comparative assessments specified in the original **C** matrix given in Eqs. 2.3 and 2.6, was then defined as

$$\mathrm{CR} = \frac{\mathrm{CI}}{\mathrm{RI}}. \tag{2.9}$$

Consistency ratios where $\mathrm{CR} \geqslant 0.1$ are generally considered to be too large to ensure reasonable consistency of the relative importances chosen in a particular case and, thus, they are usually adjusted to produce values within the range $0 < \mathrm{CR} \leqslant 0.1$. Saaty[8] has described an algorithmic approach for accomplishing this.

As will be seen in the following section, the weights associated with the relative importance of each of the ROF properties will be combined with their related utility functions to yield a total utility that lies in the unit interval [0,1] of the real line. When multiplied by the number of properties ($N_{Prop} = 4$) it will yield an ROF-Score that is identical to the number of flags set (i.e., ROF violations) in the case of classical ROF and an ROF-Score that lies on the [0,4] interval of the real line in the case of the soft ROF (vide infra).

2.2. The Rule of Five (ROF) redux

The new formulation of the ROF based on the AHP will allow generalization of the classical ROF to a new, softer version that provides a more realistic assessment of the contribution of each of the four properties to the final ROF-Score.

Consider the set of ROF properties P and the set of molecules M given in Eqs. 2.1 and 2.2, respectively. For a specific molecule M_i each property P_π defines a *utility function* u_π that maps the property value p_π, which belongs to the value set V_π associated with that property, into the set of binary integers in the classical ROF case,

$$u_\pi : V_\pi \rightarrow \{0, 1\} \tag{2.10}$$

and onto the unit interval of the real line [0,1]

$$\underset{\sim}{u}_\pi : V_\pi \rightarrow [0, 1] \tag{2.11}$$

in the 'soft' ROF case. The tilde '$\sim$' under the symbol 'u' for the soft utility function in Eq. 2.11 is used to differentiate it from the classical 'hard' utility function in Eq. 2.10. The soft utility function $\underset{\sim}{u}$ is continuous for MW and Log$P$ and discrete (non-negative integer) for HBD and HBA. A utility function value of '0' indicates that the associated property value falls within a desirable range, while a value of '1' indicates that the property value falls within an undesirable range. As shown in Eq. 2.11, soft utility functions also have values that lie within the interval (0,1). Such values can be interpreted as the 'degree of undesirability' associated with the value of a given property. Figures 2 and 3, which are discussed in Sections 3 and 4, respectively, provide graphical examples.

The values produced by either of the utility functions are weighted and summed over the four properties to give the total utilities for the classical and soft ROF methods, respectively,

$$u_{TOTAL} = w_{MW} \cdot u_{MW} + w_{LogP} \cdot u_{LogP} + w_{HBD} \cdot u_{HBD} + w_{HBA} \cdot u_{HBA} \tag{2.12}$$

$$\underset{\sim}{u}_{TOTAL} = w_{MW} \cdot \underset{\sim}{u}_{MW} + w_{LogP} \cdot \underset{\sim}{u}_{LogP} + w_{HBD} \cdot \underset{\sim}{u}_{HBD} + w_{HBA} \cdot \underset{\sim}{u}_{HBA} \tag{2.13}$$

Sections 3 and 4 provide detailed descriptions of the utility functions used in this work.

From Eqs. 2.12 and 2.13 it is clear that the same weights apply to both utility functions. In addition, since the weights sum to unity (see Eq. 2.7), it can be shown[13] that

$$0 \leqslant u_{TOTAL}, \underset{\sim}{u}_{TOTAL} \leqslant 1. \tag{2.14}$$

Figure 1 provides a graphical portrayal of the procedure for computing the value of the total utility function. Multiplying either u_{TOTAL} or $\underset{\sim}{u}_{TOTAL}$ by N_{Prop} yields the classical and soft ROF-Scores, respectively,

$$\text{ROF-Score} = N_{Prop} \cdot u_{TOTAL} \tag{2.15}$$

$$\text{sROF-Score} = N_{Prop} \cdot \underset{\sim}{u}_{TOTAL} \tag{2.16}$$

From Eqs. 2.14, 2.15, and 2.16 it then follows that

$$0 \leqslant \text{ROF-Score}, \text{sROF-Score} \leqslant N_{Prop}, \tag{2.17}$$

Thus, while in general ROF-Score $\neq$ sROF-Score, both scores are bounded by zero and N_{Prop}.

3. 'Classical' Rule of Five (ROF)

The usual formulation of the ROF assumes sharp property thresholds, which are exemplified by the four utility functions shown in Figure 2. Using these functions with the AHP, it is possible to reformulate the classical ROF in a way that allows for its generalization.

Consider the property values given in the second column of Table 1 for the hypothetical molecule M_1. In this example, as seen in the last column of Table 1, the property weights are taken to be equal:

$$w_{MW} = w_{LogP} = w_{HBD} = w_{HBA} = 0.25. \tag{3.1}$$

This is *implicitly* assumed in the classical ROF that treats all four properties as being equal contributors to the ROF-Score. Substituting the property values into the appropriate utility functions depicted in Figure 2 yields the values given in the third column of Table 1. Combining the weights and utility values in Eq. 2.12 yields

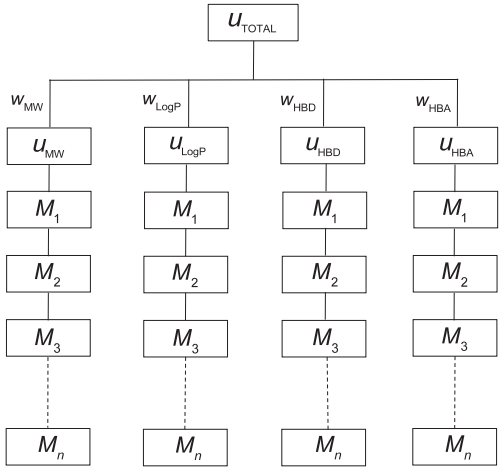

Figure 1. Graphical depiction of the three-tiered hierarchy used in the calculation of ROF-Scores (see text for further details).

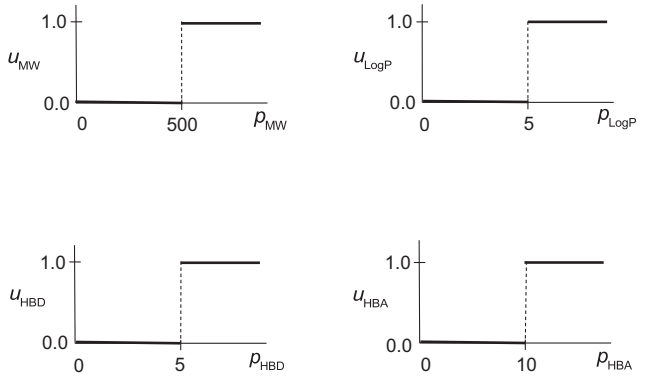

Figure 2. Utility functions with sharp boundaries ('step' functions) for the four ROF properties used to calculate ROF-Scores.

Table 1
Classical ROF for the hypothetical molecule M_1: Properties, P_π; property values, p_π; utility functions with sharp thresholds, u_π; and equal weights, w_π

P_π	p_π	u_π	w_π
MW	510	1.00	0.25
Log*P*	5.02	1.00	0.25
HBD	4	0.00	0.25
HBA	9	0.00	0.25

Table 2
Classical ROF for the hypothetical molecule M_2: Properties, P_π; property values, p_π; utility functions with sharp thresholds, u_π; and weights, w_π

P_π	p_π	u_π	w_π
MW	495	0.00	0.25
Log*P*	4.85	0.00	0.25
HBD	5	0.00	0.25
HBA	10	0.00	0.25

$$u_{\mathrm{TOTAL}}(M_1) = 0.25\,(1.00) + 0.25\,(1.00) + 0.25\,(0.00) + 0.25\,(0.00) = 0.50 \tag{3.2}$$

which upon substitution in Eq. 2.15 gives

$$\text{ROF-Score}\,(M_1) = N_{\mathrm{Prop}} \cdot u_{\mathrm{TOTAL}}(M_1) = 4\,(0.50) = 2.00. \tag{3.3}$$

Thus, two ROF flags (violations) are set for M_1.

Now consider a second hypothetical molecule M_2 whose properties are given in Table 2. Examination of the second column of Table 2 shows that M_2 has very similar properties compared to M_1. For example, the difference in molecular weight of the two molecules is only 15 Da, and their corresponding Log*P*s differ by only 0.17. Thus, it is reasonable to expect both molecules to have identical or at least similar ROF-Scores. However, the molecular weight and Log*P* of M_1 both lie above their respective thresholds, while those of M_2 lie below them. Substituting the appropriate values for M_2 from Table 2 into Eq. 2.12 yields

$$u_{\mathrm{TOTAL}}(M_2) = 0.25\,(0.00) + 0.25\,(0.00) + 0.25\,(0.00) + 0.25\,(0.00) = 0.00 \tag{3.4}$$

with a corresponding ROF-Score of

$$\text{ROF-Score}\,(M_2) = N_{\mathrm{Prop}} \cdot u_{\mathrm{TOTAL}}(M_2) = 4\,(0.00) = 0.00 \tag{3.5}$$

Thus, no ROF flags are set for M_2. This result suggests that M_1 may not be a suitable candidate for further drug research, while M_2 with very similar property values, does not carry such potential liabilities.

This surprising result lies at the heart of the major deficiency of the ROF as it is currently implemented. Fortunately, removing this deficiency is not too difficult and requires only that the current formulation be modified in a way that 'softens' the sharp thresholds used in the classical ROF.

4. Softening the Rule of Five (sROF)

4.1. 'Soft' utility functions

The first step in modifying the classical ROF is to soften the crisp boundaries used in the current implementation keeping the weights unchanged from their values in the previous example. How this can be accomplished is seen in Figure 3. In contrast to the sharp thresholds depicted in Figure 2, the thresholds in Figure 3 are distinctly different in that they are 'sloped' rather than vertical as they are in Figure 2.

It is important to note, however, that many different functional forms are possible. The ones chosen here are taken mainly for illustrative purposes and for their simplicity, an instance to be sure of Occam's razor.[14] Although no attempt was made to determine optimum utility functions for the simple functional forms used here, it will be clear from the sequel that they provide an improved account of the properties and their corresponding contributions to the ROF-Score.

As is seen in Figure 3, the utility functions corresponding to MW, Log*P*, HBA, and HBD are *piecewise linear functions* of the following general form

$$\begin{aligned} \underset{\sim}{u}_\pi &= 0 \text{ for } -\infty < p_\pi < p_\pi^{lower} \\ \underset{\sim}{u}_\pi &= \frac{p_\pi - p_\pi^{lower}}{p_\pi^{upper} - p_\pi^{lower}} \text{ for } p_\pi^{lower} < p_\pi < p_\pi^{upper} \\ \underset{\sim}{u}_\pi &= 1 \text{ for } p_\pi^{upper} < p_\pi < \infty \end{aligned} \tag{4.1}$$

where the tilde '$\sim$' under the utility function designates it as a soft utility function without a sharp threshold as in the classical case: $p_\pi = p_\pi^{\mathrm{lower}}$ is the last point in the interval $-\infty < p_\pi < p_\pi^{\mathrm{lower}}$ where the utility function has a value of zero; $p_\pi = p_\pi^{\mathrm{upper}}$ is the first point in the interval $p_\pi^{\mathrm{upper}} \leqslant p_\pi < \infty$ where the value of the utility function is one. The slope of the utility function is zero in these two intervals. The values for p_π^{lower} and p_π^{upper} are determined by the nature of the utility function, as is shown in Figure 3 for the properties MW, Log*P*, HBA, and HBD, respectively. In the region containing these two points the slope of the utility function is positive, and its value goes from $0^+ \rightarrow 1^-$. Note that the utility functions corresponding to HBA and HBD are actually discrete, although they are

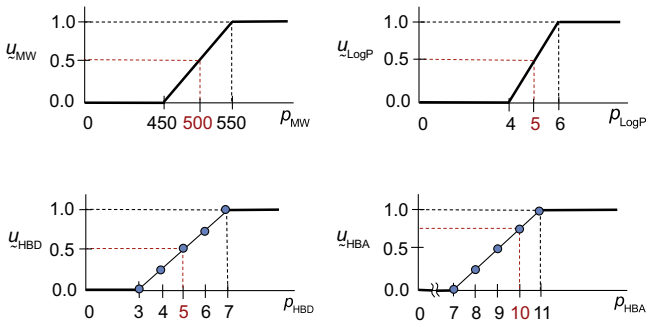

Figure 3. Utility functions with soft boundaries (piecewise linear functions) for the four ROF properties used to calculate sROF-Scores. The red numbers indicate the classical threshold values of the four ROF properties, and the dashed red lines also indicate the corresponding values of the utility functions with respect to the threshold values.

represented by continuous functions. As indicated by the filled circles that lie within the $p_\pi^{\text{lower}} < p_\pi < p_\pi^{\text{upper}}$ interval.

4.2. Equal weights

First, consider the simplest case of the soft ROF (sROF) where, as in the classical case, the weights of each of the four properties are the same as those given in Eq. 3.1. The first four columns of Tables 3 and 4 contain similar data as that given in Tables 1 and 2, except that the values of the utility functions correspond to those obtained from the soft utility functions shown in Figure 3 and Eq. 4.1. Substituting the appropriate values from Tables 3 and 4 into Eq. 2.13 yields for M_1 and M_2, respectively:

$$\underset{\sim}{u}_{\text{TOTAL}}(M_1) = 0.25(0.60) + 0.25(0.50) + 0.25(0.25) + 0.25(0.50) = 0.46 \tag{4.2}$$

and

$$\underset{\sim}{u}_{\text{TOTAL}}(M_2) = 0.25(0.45) + 0.25(0.43) + 0.25(0.50) + 0.25(0.75) = 0.53 \tag{4.3}$$

which give upon substitution into Eq. 2.16 the respective sROF-Scores:

$$\begin{aligned} \text{sROF-Score}(M_1) &= N_{\text{Prop}} \cdot \underset{\sim}{u}_{\text{TOTAL}}(M_1) = 4(0.46) = 1.84 \\ \text{sROF-Score}(M_2) &= N_{\text{Prop}} \cdot \underset{\sim}{u}_{\text{TOTAL}}(M_2) = 4(0.53) = 2.12 \end{aligned} \tag{4.4}$$

Figure 4 shows the difference between the classical and soft ROF approaches. In the classic case, M_1 exhibits two ROF violations while in the soft case both molecules exhibit very similar ROF-Scores—a much more realistic measure given the similarity of the property values for these two molecules. Although there is a transposition in the ordering from their classic ROF-Scores, the two sROF-Scores are close enough that such an interchange is not likely to be highly significant.

Clearly, the soft approach produces ROF-Scores that are in better accord with the similar property values of both molecules. However, the interpretation traditionally employed in ROF analysis, namely, that two ROF flags (violations) 'suggest' that the molecules may not be suitable for drug development must be reconsidered (vide infra).

4.3. Modifying the weights

As a final example, the utility functions for the two hypothetical molecules remain unchanged, but the weights are altered to illustrate the effect of dropping the assumption that all properties are equally important. In a system as simple as the one considered in this work, it may be possible to choose an appropriate set of weights directly. However, in many cases, especially those involving more than the four ROF properties (e.g., number of rotatable bonds, polar surface area, and molecular complexity) it is not straightforward to generate appropriate sets of weights.

Table 3
Soft ROF for hypothetical molecule M_1: Properties, P_π; property values, p_π; soft utility functions, $\underset{\sim}{u}$; and weights, w_π and w'_π

P_π	p_π	$\underset{\sim}{u}_\pi$	w_π	w'_π
MW	510	0.60	0.25	0.528
Log$P$	5.02	0.50	0.25	0.300
HBD	4	0.25	0.25	0.087
HBA	9	0.50	0.25	0.085

Table 4
Soft ROF for hypothetical molecule M_2: Properties, P_π; property values, p_π; soft utility functions, $\underset{\sim}{u}$; and weights, w_π and w'_π

P_π	p_π	$\underset{\sim}{u}_\pi$	w_π	w'_π
MW	495	0.45	0.25	0.528
Log$P$	4.85	0.43	0.25	0.300
HBD	5	0.50	0.25	0.087
HBA	10	0.75	0.25	0.085

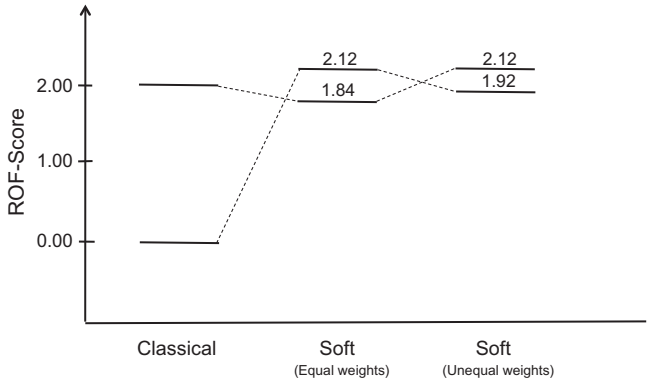

Figure 4. Comparison of classical ROF-Scores to soft ROF-Scores with equal and unequal weightings.

In this section, the AHP is employed as a means for generating consistent sets of modified weights. Consider, as an example, the comparison matrix

$$\mathbf{C} = \begin{array}{c|cccc} & \text{MW} & \text{LogP} & \text{HBD} & \text{HBA} \\ \text{MW} & 1 & 3 & 4 & 5 \\ \text{Log D} & \frac{1}{3} & 1 & 5 & 4 \\ \text{HBD} & \frac{1}{4} & \frac{1}{5} & 1 & 1 \\ \text{HBA} & \frac{1}{5} & \frac{1}{4} & 1 & 1 \end{array}. \tag{4.5}$$

The maximum eigenvalue and corresponding normalized maximum eigenvector, whose elements are the corresponding weights (see Eq. 2.7), are given by

$$\lambda_{\max} = 4.168, \quad \mathbf{v}^{\max} = \begin{bmatrix} 0.528 \\ 0.300 \\ 0.087 \\ 0.085 \end{bmatrix} \begin{matrix} \text{MW} \\ \text{LogP} \\ \text{HBD} \\ \text{HBA} \end{matrix}. \tag{4.6}$$

With respect to pairwise comparisons, the elements in the first row of the matrix in Eq. 4.5 show that MW is three times as important as Log$P$, four times as important as HBD, and five times as important as HBA. Because of the reciprocal nature of the matrix, the elements of the first column show that Log$P$ is one-third as important, HBD is one-quarter as important, and HBA is one-fifth as important as MW. *This example is only intended to be an illustrative not a definitive example of the weighting procedure*. The normalized eigenvector (weight vector) given in Eq. 4.6 shows that MW and Log$P$ are the most important factors (in this example, not necessarily in general), while HBA and HBD have a diminished effect.

Although it is not discussed here, *sensitivity analysis* can provide a means for assessing the influence that perturbations of the elements of the C matrix can have on its maximum eigenvalue and eigenvector.[8] The value of the consistency index given in Eq. 2.8, and the corresponding random consistency index are equal to 0.056 and 0.90, respectively. Thus, the consistency ratio, CR = CI/RI given in Eq. 2.9

is equal to 0.062, which lies below the consistency threshold of 0.1 typically used in this method. It is, however, important to note that consistency is not, of itself, a measure of how relevant the chosen values of the pairwise comparisons are to the problem, only that the values chosen are to a large degree internally consistent.

Returning to Tables 3 and 4, combining the values for the soft utility functions, which remain unchanged from the previous case, with the modified weights, w'_π, estimated by the AHP procedure given in the fifth columns of the tables yields the new total utilities for M_1 and M_2

$$\begin{aligned} \underset{\sim}{u}_{\mathrm{TOTAL}}(M_1)' &= 0.528(0.60) + 0.300(0.50) + 0.087(0.25) \\ &\quad + 0.085(0.50) \\ &= 0.53 \end{aligned} \tag{4.7}$$

and

$$\begin{aligned} \underset{\sim}{u}_{\mathrm{TOTAL}}(M_2)' &= 0.528(0.45) + 0.300(0.43) + 0.087(0.50) \\ &\quad + 0.085(0.75) \\ &= 0.48 \end{aligned} \tag{4.8}$$

which upon substitution into Eq. 2.16 yield

$$\text{sROF-Score}(M_1)' = N_{\mathrm{Prop}}\underset{\sim}{u}_{\mathrm{TOTAL}}(M_1)' = 4(0.53) = 2.12 \tag{4.9}$$

and

$$\text{sROF-Score}(M_2)' = N_{\mathrm{Prop}}\underset{\sim}{u}_{\mathrm{TOTAL}}(M_2)' = 4(0.48) = 1.92, \tag{4.10}$$

respectively.

Figure 4 provides an overall picture of the results for the sharp and soft threshold approaches to the computation of ROF-Scores for all three cases. From the figure it is clear that both sROF calculations provide similar results. However, the ordering is different in the two cases, as there is a crossover between the soft values obtained with and without equal weights. Nonetheless, both of the soft results are in better accord with the property values of the two hypothetical molecules than is the classical, sharp threshold result, where for M_1, two ROF flags are set while for M_2 no flags are set.

Such crossovers, which are discussed further in Section 5.2, are reasonably common for both equal and unequal weights. No attempt was made in this work, however, to determine optimal weightings, although it is clear that the choice of weightings can markedly affect the final orderings obtained (vide infra). Determination of such optimal weightings will be dealt with in future work.

5. Results and discussion

5.1. DrugBank Library (DBL)

In order to better appreciate how this methodology performs on real datasets, a sample of 1044 compounds was extracted from the DBL,[15] which consists of more than 4000 drugs and related compounds. The sample was obtained as follows. First, the DBL sublibraries 'Approved Structures' and 'Small Molecules' were intersected yielding a library of 1356 so-called 'Approved Small Molecules' (as of the June 22, 2010 version of the database). Next, compounds without structure or with an 'Illicit Drug' annotation were removed. Finally, salts were stripped off and compounds with molecular weights in the range between 200 and 700 Da were retained, yielding a final set of 1044 approved small molecules. An Excel file containing the identifier numbers for the subset of DBL compounds used in this study is available electronically as Supplementary data. The sample set was imported into MOE (Chemical Computing Group, Montreal)[16] in sdfile format, and the compounds of the resulting dataset were characterized by the four ROF properties, which are based on the MOE descriptors *Weight*,

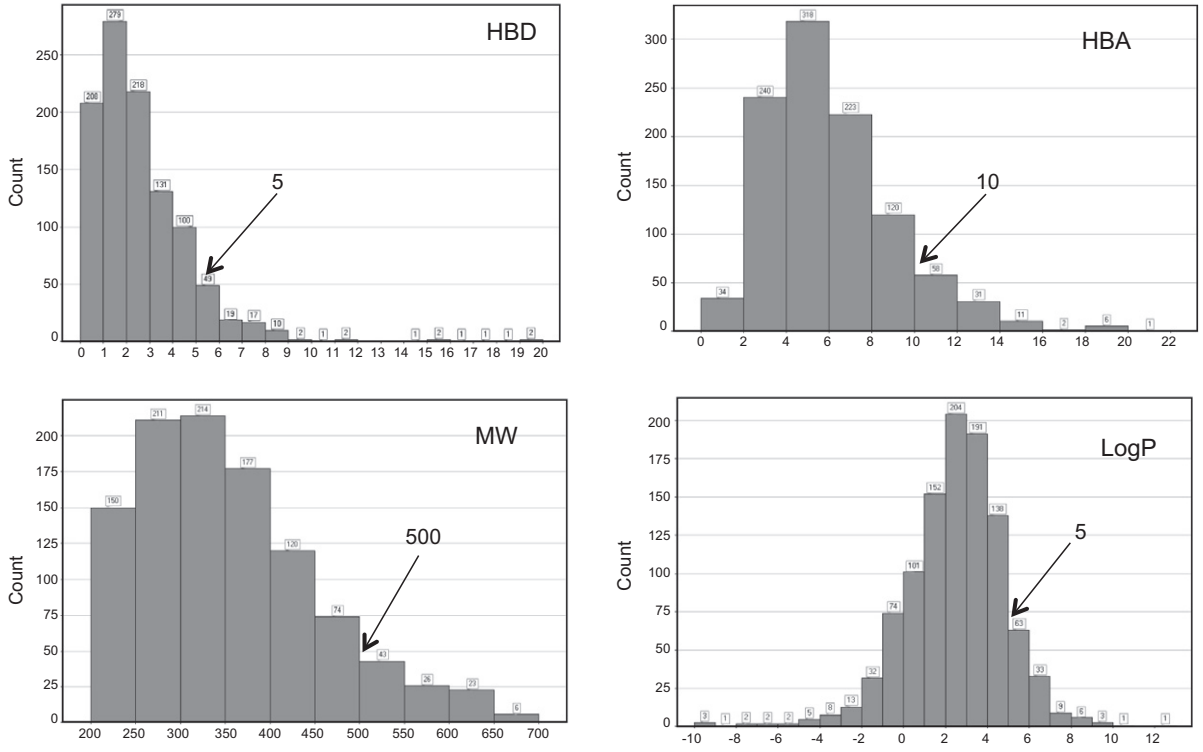

Figure 5. Histograms showing the distribution of property values for the sample of 1044 DBL Compounds considered in this work. The vertical arrows point to the approximate threshold values used for each of the four ROF properties.

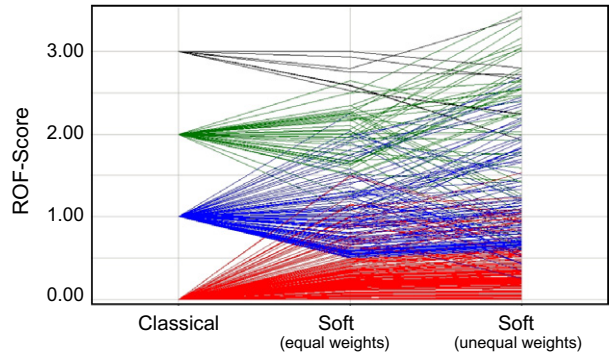

Figure 6. Computed ROF-Scores for the sample of 1044 DBL compounds considered in this work. The figure illustrates the significant amount of crossover for both of the sROF-Scores (equal and unequal weights) compared to the corresponding classical ROF-Score.

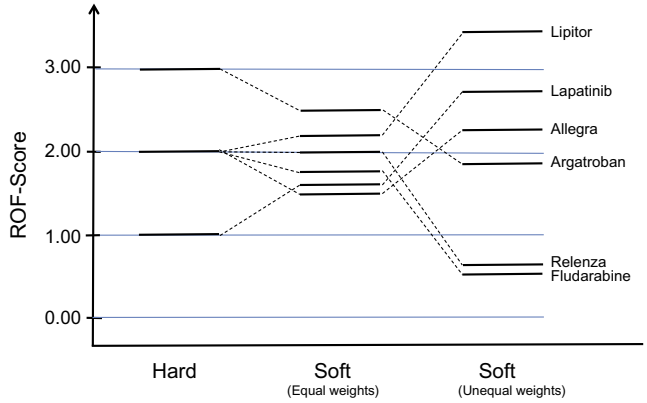

Figure 7. Diagram depicting the ROF- and sROF-Scores for six drug molecules taken from the DBL.

LogP(o/w), *lip_acc*, and *lip_don* for the properties MW, Log*P*, HBA, and HBD, respectively.

Figure 5 provides a summary of the distributions of the four chemical properties employed in the ROF. The arrows in the figure point to the *sharp thresholds* for the four rules used in the classical ROF computations. From the figure it is apparent that a reasonable fraction of the compounds violate at least one of the four rules, except perhaps for the rule associated with Log*P*. Thus, they provide a reasonable basis for the analysis of the soft ROF procedure that follows.

5.2. Soft ROF analysis of the DBL sample

Applying the computational procedures described above to the set of 1044 DBL compounds yields the plot shown in Figure 6, which shows that in the case of classical ROF-Scores the number of compounds decreases with respect to the number of rule violations, although none of the compounds has four ROF violations.

The situation is distinctly different in the case of sROF-Scores where there is considerable spread in both the equally and unequally weighted scores. In comparison to their classical counterparts a number of the equally weighted sROF-Scores exhibit crossover. In crossovers some of the compounds with lower classical ROF-Scores have higher sROF-Scores than some of the compounds with higher classical ROF-Scores. The most prevalent crossovers occur between compounds with respective ROF-Scores of 0 and 1, with somewhat fewer crossovers between compounds with ROF-Scores of 1 and 2. Only two crossovers occur between compounds with ROF-Scores of 2 and 3. Crossovers are also observed between the equally- and unequally-weighted sROF-Scores, indicating that the results obtained are, not surprisingly, dependent on the weightings chosen. How to optimally choose these weightings is, however, an area of future research. The current example of unequally weighted sROF-Scores is provided merely to illustrate that the weightings can be used to emphasize the contribution of particular properties to sROF-Scores.

A more detailed analysis of six drugs found in the DBL sample provides additional insights into the workings of the utility function approach for determining ROF-Scores. Details given in Table 5 and Figure 7 provide a summary of the three types of ROF-Scores described. As shown in the table and the figure, Lapatinib is the only compound with a ROF-Score = 1.00, that is it has only a single classical ROF violation, which is due to its MW of 581.1 Da. Figure 3 shows that its MW is high enough for its utility function to have a value of unity. On the other hand, its log*P* and HBA values, 4.84 and 8, respectively, are both less than their respective ROF thresholds, so no classical ROF flags are set. They do, however, lie close enough to their threshold values to contribute to the sROF-Score as shown in Figure 3. Thus, Figure 7 shows that Lapatinib's sROF-Score increases over its corresponding ROF-Score. Now consider Argatroban, which has three ROF violations associated with its MW of 508.6 Da, its HBD of 7, and its HBA of 11. Again, Figure 3 shows that its MW has a utility-function value of about 0.5, while the corresponding values for HBD and HBA are both unity. Thus, the lower value associated with MW yields an overall sROF-Score that falls below its ROF-Score of 3.00. The sROF-Scores of the remaining five compounds lie close to their corresponding classical ROF-Scores of 2.00.

Now consider the case where the weights are modified, as shown in Eqs. (4.7)–(4.9). The modified weightings clearly emphasize the contributions of MW and Log*P* at the expense of HBA and HBD. Note that it is not claimed that this relationship is necessarily true, only that it provides an example. As seen in Figure 7, the most interesting cases are for Lipitor, Relenza, and Fludarabine where the differential weightings have a considerable impact on their sROF-Scores. In the case of Lipitor, which has a MW of 558.7 Da and a Log*P* of 6.47, its utility-function value is equal to unity. Since these two values are heavily weighted (0.528 and 0.300, respectively), they will have a strong impact on the total utility value and, thus, on the overall sROF-Score. Relenza and Fludarabine, on the other hand, have MW's and Log*P*'s whose utility-function values are zero, so the heavy weighting on these two properties has no effect on their final sROF-Scores. Although the membership function values for HBD and HBA are close to unity their associated weights of 0.087 and 0.085, respectively, are insufficient to raise their sROF-Scores above their ROF-Scores of 2.00.

These examples are meant to illustrate how the softening of the classical ROF provides a more reasonable scoring system that places molecules with similar property values close to one another, which is not always the case with the classical ROF. And, if desired, the use of unequal weightings can provide 'adjustments' to the utility-function values that take account of differing importance of their properties. As can be seen from the data in Table 5 and Figure 7, the differential weightings can have a significant effect on the sROF-Scores obtained. This may be of especial importance in extended ROF models that take account of additional properties such as the number of rotatable bonds or the polar surface area, to name two of the many possibilities. Lastly, the differential weighting scheme used here can also be applied to the classical ROF method, although to our knowledge it has not been.

5.3. Utility functions and weights

Certainly the choice of functional form for the utility functions has an impact on the final results. However, the preliminary work

Table 5
Summary of the properties of six drugs taken from the DBL, along with their corresponding classical and soft ROF-Scores

Drug Bank ID	Drug name	Molecular properties[a]				ROF-Scores		
		MW (Da)	Log*P*	HBD	HBA	Class.[b]	Soft[b]	Soft[c] (wt'd)
DB00278	Argatroban	**508.6**	1.62	**7**	**11**	3	2.57	1.93
DB00558	Relenza	332.3	−3.13	**9**	**11**	2	2.00	0.69
DB00950	Allegra	**501.7**	**7.30**	3	5	2	1.52	2.29
DB01073	Fludarabine	365.2	−2.69	**6**	**12**	2	1.75	0.60
DB01076	Lipitor	**558.7**	**6.47**	4	7	2	2.25	3.40
DB01259	Lapatinib	**581.1**	4.84	2	8	1	1.67	2.70

[a] Property values given in boldface type exceed classical ROF thresholds.
[b] Utility function values are weighted equally.
[c] Utility function values are differentially weighted.

described here shows that softening the computation of ROF-Scores tends to spread the molecules along the 'score axis' as depicted in Figures 6 and 7. This provides a more fine-grained representation that is less sensitive to small differences in property values, especially in comparison to the usual implementation of the classical ROF. Compounds can now be chosen based on some pre-determined cutoff value or they can be selected to lie within a given range of sROF-Scores. Although not investigated in this work, sensitivity analysis is a reasonable next step in the development of this methodology (see e.g., Ref. 8)

The fact that there are numerous crossovers observed in the sROF-Scores with respect to the classical ROF-Scores also shows the drastic effect that the severe discretization embodied in the ROF approach produces. Finally, the crossovers observed between the equally and unequally weighted sROF-Scores clearly indicates the need for further work to determine more optimum weightings, a subject that will be dealt with in future work.

5.4. Is model validation appropriate here?

Model validation is a requirement in the development of a majority of cheminformatics and modeling methods. However, there are cases where it may not be entirely appropriate. ROF- and sROF-Scores also fall into this category and are based upon a statistical analysis of the properties of compounds from a very large dataset.[7] The number of property values that exceed their given thresholds and the relationship of the corresponding ROF-Score to any global characteristic such as 'drug-likeness,' which is an ill-defined concept, is tenuous at best. And, as is the case for similarity, strict validation of the methodology is not possible because there are no completely objective values that can be used as a basis for testing. Nevertheless, the ROF, like similarity, has been quite useful in many practical applications. This begs the question of how to evaluate the sROF method proposed here. As has been shown earlier, the sROF produces linearly ordered lists of compounds that are more in accord with their relative property values. Compounds with similar values tend to lie near one another in a given list, although whether the rankings are related to drug-likeness remains problematic. The main importance of the sROF, in contrast to the classical ROF, is its robustness compared to the latter. As illustrated by the work presented here, small differences in property values do not result in large differences in ROF-Scores.

Thus, since the classical ROF has been shown to be useful in many applications, it is not unreasonable to assume that the sROF will be at least as useful as its parent.

6. Summary and conclusions

A new approach, based on decision theory, namely, the AHP, has been developed that extends the current classical ROF by removing the sharp thresholds for the four properties (MW, Log*P*, HBA, and HBD) that are the basis of the method. In the new, softer approach, compounds that have similar property values now tend to lie near one another in the 'space' of sROF-Scores. This is clearly a more desirable situation than that provided by the classical ROF formulation where compounds with similar property values may have significantly different ROF-Scores. Softening the ROF procedure is based on replacing the sharp ('step') utility function of the classical ROF, which takes on values of either '0' or '1', by softer versions, which take on values over unit interval [0,1] of the real line.

In addition to softening of the thresholds for the four ROF properties, their relative importance can be accounted for by differentially weighting their values. In simple cases, such as illustrated in this work, weightings can be determined *ad hoc*, but this may become more difficult if additional properties are considered. In such cases, the AHP described in this work can provide a suitable means for determining a reasonable set of weights. It is also possible to develop much more complex, multi-layered AHP-based decision schemes than the simple ones presented here (see e.g., Ref. 8,10). Importantly, such an enhanced approach holds promise as an alternative means for developing soft decision-theoretic models in lieu of the more rigid mathematical models currently used for assessing the drug-likeness of compounds.

The ROF has become a de facto standard for assessing drug-like properties of molecules (vide supra) in compound collections, in compound acquisition, and in combinatorial library design to name a few of its applications. Recent work by Burton, Petit, Meurice, and Maggiora on compound acquisition based on the soft ROF approach showed that it provided a more robust and consistent characterization than was obtained using the traditional ROF approach [Unpublished work, 2011]. Although this is only a single example, it provides additional evidence that suggests the soft ROF approach may afford a better, albeit imperfect, description of 'drug-like' properties than the traditional ROF procedure in widespread use today.

The utility function procedure described here is presented as an alternative to the present ROF method, but it is not meant to be a definitive treatment of the subject. Rather, it is meant to illustrate how taking a softer approach can lead to a more robust handling of the ROF and provide results that are more in accord with expectations based upon the property values of the molecules being considered. However, details and limitations of the methodology await the results of future work.

More generally, the method described in this work illustrates how the artifacts associated with the sharp thresholds inherent in many discrete-mathematical methods may in some cases be ameliorated using the softer procedures. Thus, the soft approach can potentially be extended to many problems of interest in drug discovery and in other fields where sharp thresholds are present.

Acknowledgment

Award Number RC2MH090878 from the National Institute Of Mental Health provided partial support for the work described in this paper. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institute of Mental Health or the National Institutes of Health.

Supplementary data

Supplementary data associated with this article can be found, in the online version, at doi:10.1016/j.bmc.2011.11.064. These data include MOL files and InChiKeys of the most important compounds described in this article.

References and notes

1. Clark, D. E.; Pickett, S. D. *Drug Discovery Today* **2000**, *5*, 49.
2. Muegge, I.; Heald, S. L.; Brittelli, D. *J. Med. Chem.* **2001**, *44*, 1841.
3. Blake, J. F. *Curr. Opin. Biotechnol.* **2000**, *11*, 104.
4. Sadowski, J.; Kubinyi, H. *J. Med. Chem.* **1998**, *41*, 3325.
5. Walters, W. P.; Murcko, M. A. *Adv. Drug Deliv. Rev.* **2002**, *54*, 255.
6. Oprea, T. I.; Davis, A. M.; Teague, S. J.; Leeson, P. D. *J. Chem. Inf. Comput. Sci.* **2001**, *41*, 1308.
7. Lipinski, C. A.; Lombardo, F.; Dominy, B. W.; Feeney, P. J. *Adv. Drug Deliv. Rev.* **1997**, *23*, 3.
8. Saaty, T. L. Fundamentals of Decision Making and Priority Theory with the Analytic Hierarchy Process In *Vol. VI of the AHP Series*; RWS Publications: Pittsburg, Pennsylvania, 2000.
9. Saaty, T. L.; Vargas, L. G. Decision Making in Economic, Political, Social, and Technological Environments with the Analytic Hierarchy Process In *Vol. VII of the AHP Series*; RWS Publications: Pittsburg, Pennsylvania, 1994.
10. Maggiora, G. M. Computer-aided Decision Making in Pharmaceutical Research. In *Proceedings of the International Workshop on Molecular Informatics—Confronting Complexity*; Hicks, M. G., Kettner, C., Eds.; Beilstein Institute: Frankfurt, 2003; p 149.
11. Wager, T. T.; Hou, X.; Verhoest, P. R.; Villalobos, A. *ACS Chem. Neurosci.* **2010**, *1*, 435.
12. Minc, H. *Nonnegative Matrices*; Wiley-Interscience Publication: John Wiley & Sons, New York, 1988.
13. Torra, V.; Nurakawa, Y. *Modeling Decisions—Information Fusion and Aggregation Operators*; Springer: New York, 2007.
14. Wikipedia: http://en.wikipedia.org/wiki/Occam%27s_Razor. Accessed September 20, 2011).
15. Wishart, D. S.; Knox, C.; Guo, A. C.; Shrivastava, S.; Hassanali, M.; Stothard, P.; Chang, Z.; Woolsey, J. *Nucleic Acids Res.* **2006**, *34*, D668. (http://www.drugbank.ca/databases. Accessed 11.21.10).
16. MOE (The Molecular Operating Environment) Version 2008.10, software available from Chemical Computing Group Inc., 1010 Sherbrooke Street West, Suite 910, Montreal, Canada H3A 2R7. (http://www.chemcomp.com. Accessed 11.21.10.