Chemistry & Biodiversity (in press) doi: 10.1002/cbdv.200900148

Beyond Profiling: Using ADMET models to guide decisions

Matthew Segall, Edmund Champness, Olga Obrezanova, Chris Leeding

Optibrium Ltd, 7226 IQ Cambridge, Beach Drive, Cambridge, CB25 9TL, UK

Abstract

ADMET models, whether *in silico* or *in vitro,* are commonly used to 'profile' molecules, to identify potential liabilities or filter out molecules expected to have undesirable properties. While useful, this is the most basic application of such models.

We will demonstrate how models may be used to go 'beyond profiling' to guide key decisions in drug discovery. For example, selection of chemical series to focus resources with confidence or design of improved molecules targeting structural modifications to improve key properties.

To prioritise molecules and chemical series, the success criteria for properties and their relative importance to a project's objective must be defined. Data from models (experimental or predicted) may then be used to assess each molecule's balance of properties against those requirements. However, in order to make decisions with confidence, the uncertainties in all of the data must also be considered.

In silico models encode information regarding the relationship between molecular structure and properties. This is used to predict the property value of a novel molecule. However, further interpretation can yield information on the contributions of different groups in a molecule to the property and the sensitivity of the property to structural changes. Visualising this information can guide the redesign process.

In this paper we describe methods to achieve these goals and drive drug discovery decisions and illustrate the results with practical examples.

Introduction

The enormous cost of pharmaceutical R&D is driven by the high failure rate of molecules from the earliest phases, where thousands of molecules may be synthesized and screened in a single project, through to the most expensive clinical phases, where the success rate remains less than one in ten [1]. The reasons for these failures are numerous; however it is now widely accepted that a successful drug must satisfy a wide range of criteria, including absorption, distribution, metabolism and elimination (ADME), physicochemical and safety properties, in addition to potency against its therapeutic target.

To address this need, many approaches to measuring or predicting these properties in early drug discovery have been developed [2,3]. These allow data to be generated that provide information about critical properties on large numbers of molecules. However, this brings a new challenge: How to use this data to guide decisions that will enable resources to be quickly focussed on molecules that are likely to succeed downstream, thus improving the overall efficiency of the R&D process. This challenge, often described as 'multi-dimensional optimisation', is made even more difficult by the fact that these data sources, whether *in silico*, *in vitro*, or *in vivo*, are all models of the ultimate target system, a human patient. Furthermore, the data generated often exhibit significant uncertainty due to statistical errors in predictive models or variability in experimental assays.

This paper will not focus on the methods for generating data in drug discovery, but on approaches for using this data effectively, guiding key decisions regarding the selection and design of molecules. Specifically, we will look at approaches for identifying molecules likely to exhibit an appropriate balance of properties to meet the therapeutic objectives of a project and for using the information captured by *in silico* models to guide the design of molecules with improved properties.

We will contrast these novel approaches with simple 'profiling' that is commonly used to analyse the data for multiple properties. In profiling, the predicted or measured properties of a chemical series or molecule are compared against the project's criteria for success, or target product profile. Profiling may be a useful tool for spotting patterns in properties, such as consistent 'passes' or 'failures' against a given criterion. However, there are some significant limitations to the information that simple profiling provides. For example, a 'failure' against one criterion may be critical, while another could be less important to the outcome of the project – it may be appropriate to 'trade off' one property to achieve a better outcome in another.

In the first section we will describe simple profiling and illustrate some of its strengths and weaknesses. Following this, we will describe an approach that goes beyond this to provide a more comprehensive interpretation of the data. Finally, we will describe a novel method that gives additional interpretation of a predictive model to guide redesign of molecules to further refine their properties. These will be illustrated using practical examples from drug discovery projects and we consider some lessons learned regarding the application of these techniques in the Conclusions section.

Profiling

Profiling a chemical series or molecule involves the comparison of predicted or measured values for a number of properties against pre-determined criteria that define a desirable molecule. These criteria are often described as a target product profile (TPP) and are typically chosen based on experience from previous projects. The TPP will vary from project to project, depending on the therapeutic target, intended route of administration and, potentially, commercial requirements such as differentiation from competitor molecules. The TPP may change as the project proceeds, usually (but not always) imposing more strict criteria during later phases.

An example of a simple TPP for predicted ADME properties is shown in the first two columns of the table in Figure 1, in this case designed to identify appropriate molecules for an orally dosed molecule against a non-CNS (peripheral) target. This is representative of a typical TPP for a hit-to-lead project, where potent hit molecules have been identified, but no further *in vitro* or *in vivo* data have been generated.

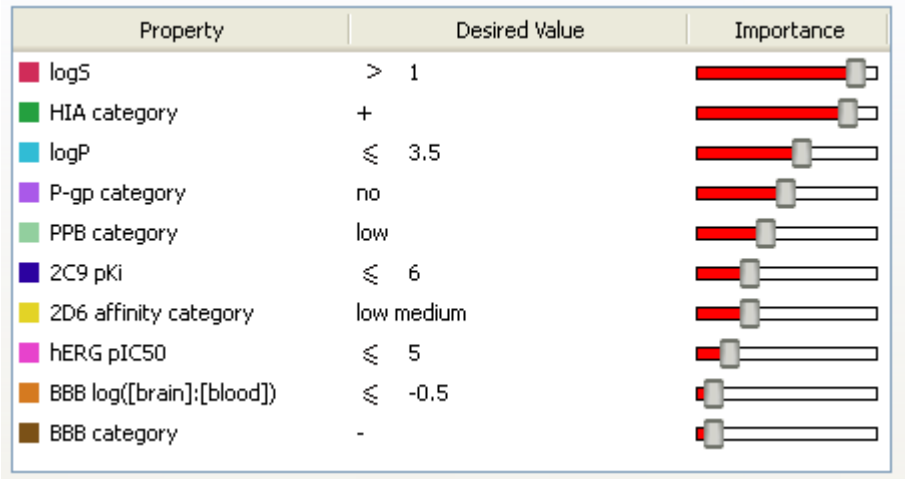

Figure 1 Example of a simple scoring profile for a range of predicted ADME properties, chosen to identify compounds suitable for an orally dosed compound for a CNS target. The first two columns represent a target product profile (TPP), but a scoring profile requires additional information on the relative importance of the criteria, as represented by the length of the red bar in the third column.

Profiles for chemical series are conveniently visualized by plotting a histogram in which each bar represents the percentage of molecules in the chemical series that meet a criterion for a single property. Examples of these are shown in the top row of Figure 2. When examining properties on a molecule-by-molecule basis, 'traffic lights' that colour code properties according to whether they pass (green) or fail (red) the corresponding criterion or are close to the criterion boundary (yellow-orange) are often used. While visually appealing, this view rapidly becomes very complicated when dealing with large numbers of properties and, if an ideal (all green) molecule is not present, it is difficult to select molecules visually; for example is it better to have one red property value or three yellow?

In order to compare the results of simple profiling with scoring methods described below, we will use an example of a hit-to-lead project, in which a diversity-based screen yielded hits from multiple chemotypes, or arrays. The challenge for this project

team was to select a smaller number of chemistries in order to focus the use of their limited synthetic and experimental resources.

The histograms in the top row of Figure 2 show the results of a simple profile, using the TPP in Figure 1, for the molecules in three arrays, labelled A, B and C. From this analysis it is immediately clear that all of the molecules in array C fail to meet the criteria for five properties; solubility, logP, plasma-protein binding, CYP2C9 affinity and hERG inhibition. This indicates that the chemistry in array C is high-risk and the likelihood of designing a molecule that would overcome all five liabilities is low. However, comparing the profiles for arrays A and B provides less information. Both arrays show a similar rate of 'failures' but these are distributed across different properties. Are the 'failures' in array A more, or less, serious than those in array B? Furthermore, is either chemistry likely to yield molecules that will pass all of the criteria simultaneously?

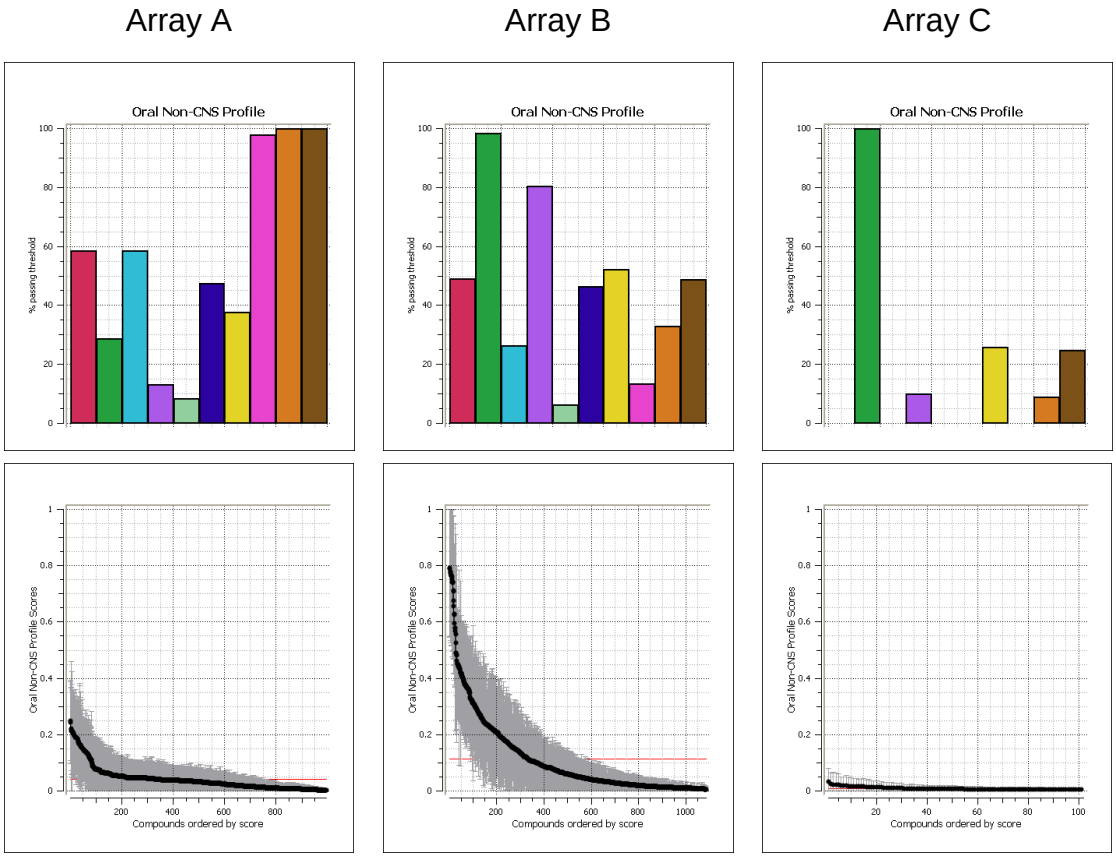

Figure 2 Profiles and scoring plots comparing three chemical arrays, A, B and C. The histograms in the top row indicate the percentage of compounds meeting the criterion for each predicted ADMET property as defined by the TPP in Figure 1. The colours of the bars correspond to the key in Figure 1. The second row shows scoring plots for the same arrays (as described in greater detail in Figure 3). Qualitatively, the greater area under the plot above for Array B indicates a higher likelihood of identifying high quality compounds in this chemical series.

Beyond Profiling - Probabilistic Scoring

A more comprehensive alternative to profiling is to generate a score that reflects the overall quality of each molecule based on the available data and their relative importance within the TPP. The methods underlying the probabilistic scoring employed herein are discussed in more detail in [4] but here will give a brief overview.

A probabilistic score is one which indicates the probability of success of a molecule against some criteria (e.g. a TPP). To score molecules for a TPP, a scoring profile must be defined, similar to a TPP described above. However, in addition to each criterion in the TPP, it is also important to specify their relative importance, as in practice it is often necessary to make a trade-off between properties if an ideal molecule cannot be identified. This is illustrated in the third column of Figure 1 for the TPP shown. Furthermore, more subtle trade-offs can be defined than simple pass/fail criteria, as a scoring profile could contain more complex functions for each property representing a range of acceptability over the property value range.

When combining data on multiple properties, it is also important to consider the uncertainty in each data point, as this could lead to the overall uncertainty in the scores being high, reducing our ability to confidently distinguish high and low quality molecules.

The result of this process is a score for each molecule, representing the likelihood of a molecule meeting the scoring criteria and an uncertainty in the overall score, derived from the uncertainties in each of the individual property values. These uncertainties can be used to establish whether the available data allow molecules to be confidently distinguished, i.e. when one molecule can be confidently chosen over another. An illustration of the output for a small set of molecules is shown in Figure 3.

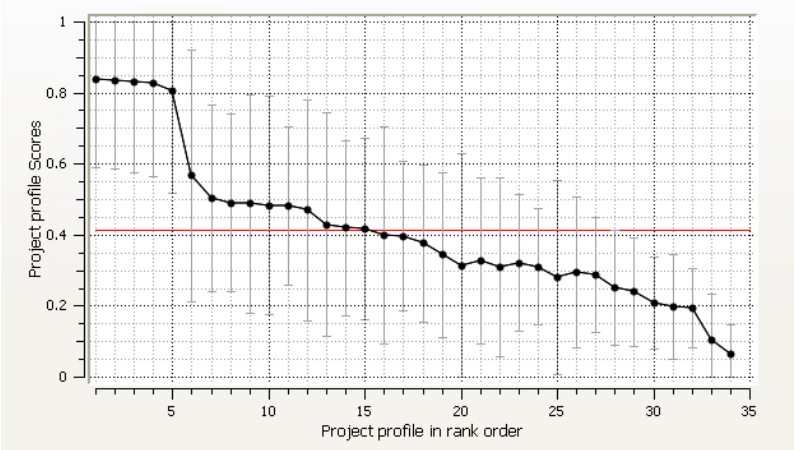

Figure 3 In this graph, molecules are plotted along the *x*-axis in rank order. The score is plotted on the *y*-axis, with error bars indicating the overall uncertainty in the score. Here the top 5 compounds cannot be confidently distinguished; more data or further criteria are required to choose between these. However, ~50% of compounds are significantly less likely to meet the project criteria than the top 5.

Scoring the arrays illustrated in Figure 2 provides significant additional information, as shown by scoring plots in the bottom row of this figure. It remains easy to see that array C corresponds to high risk chemistry, the chances of success of these molecules are very small and there is a high degree of confidence in this assessment. However, there is now a clear difference between arrays A and B. It can be clearly seen that there are a number of molecules that have a high chance of success, i.e. of meeting the scoring criteria. Conversely, the overall chances of success of the best molecules in array A are significantly lower. Therefore, one can conclude that the majority of resources can confidently be focussed on array B. Due to the statistical uncertainties in the scores, it may also be valuable to expend a small effort on the highest scoring molecules in array A to confirm this conclusion.

Based on this analysis, the molecules from hit series B were selected for resynthesis and experimental study. A range of *in vitro* ADME properties were measured for these molecules which were re-scored based only on their *in vitro* properties. The results confirmed the predicted hypothesis and allowed efforts to be further focused and rapid progress to be made.

Guiding Redesign - Glowing Molecule

As illustrated above, predictive models may be very effective for selecting molecules or chemical series when used in an appropriate framework. However, when looking in detail at smaller numbers of molecules a common criticism is that models give no indication as to why a molecule is predicted to have a certain property value, or how a molecule may be improved. Models encode relationships between molecular structure and properties, but interpreting and visualising this information to design better molecules has been almost impossible. This is particularly true of models built with modern 'machine learning' techniques such as artificial neural networks [5], Gaussian processes [6,7] or support-vector machines [8]. The models that these techniques create have commonly been described as 'black box.'

The Glowing Molecule method analyses the mathematical relationships captured by predictive models to create an intuitive visualisation of the impact of structural features of a molecule on a predicted property. The output of the analysis is a coloured field on which the 2D molecular structure is superimposed which clearly identifies 'problem' regions on a molecule and highlights functional groups that tend to improve a predicted molecular property. This information provides a guide to the changes that are most likely to result in a molecule with improved properties.

The underlying method is based on consideration of a predictive model as a mathematical function that relates a set of descriptors (x_1,x_2,x_3…) that characterise a molecule to a value of the property being modelled (y). Common descriptors include; simple 1D descriptors (e.g. molecular weight or atom counts), 2D descriptors (e.g. molecular fragments or topological polar-surface area), 3D descriptors that capture information about the shape of a molecule, and whole-molecule properties such as logP. The mathematical function f(x_1,x_2,x_3…) that correlates with the predicted property is typically fitted to a data set of molecules with known property values, using a statistical or 'machine learning' technique.

The mathematical function f(x_1,x_2,x_3…) can be considered to define a hypersurface in the space of descriptors. In the simplest, linear case, f(x_1,x_2,x_3…)=c_1x_1+c_2x_2+c_3x_3…,

this is simply a plane. However, for non-linear models such as those created using more advanced techniques, this hypersurface can adopt more complex forms. The act of making a prediction for a molecule can be represented as finding the 'height' (y) of this hypersurface for a given set of coordinates (x_1,x_2,x_3…). However, the 'shape' of this surface contains additional information, in particular the descriptors that contribute most to variation in the property value.

It is notable that these trends are constant for a linear model; they do not depend on the particular molecule for which the prediction is being made. This reflects the relatively straightforward interpretation of linear models, where the influence of descriptors is simply related to the magnitude and sign of the coefficients in the equation (c_1,c_2,c_3…) [9]. However, the trends may vary significantly between molecules for non-linear forms of f.

If the contributions of each atom in a molecule to each descriptor can also be quantified, this enables the contributions of each atom to the overall trend in the molecule's property to be calculated. This is most easily achieved for 1D and 2D descriptors, but may be generalised to more complex descriptors.

In order to illustrate the application of Glowing Molecule, we will use an example from the public domain, in order to avoid issues of confidentiality regarding molecule structures. Rowley et al. [10] investigated a series of tryptamines with very high affinity and selectivity for the h5-HT2A receptor. Their main issue was to reduce affinity for the hERG IKr channel, measured by displacement of 4 nM [3H]-dofetilide binding to HEK cells stably expressing the hERG channel.

We have applied the Glowing Molecule analysis to this series, using the QSAR model of hERG pIC50 available within the StarDrop™ software package. As this model is based on 'gold standard' patch-clamp measurements, we expect only qualitative agreement with the dofetilide displacement results reported. Despite this, Figure 4 demonstrates that the resulting visualisation provides good guidance on structural modifications.

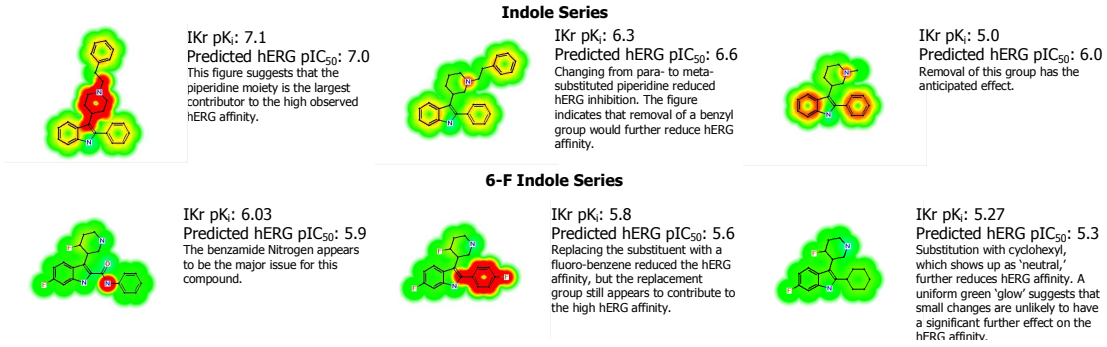

Figure 4 Examples of the Glowing Molecule visualisation of hERG inhibition for two series of hERG inhibitors from [9]. The pKi for hERG channel inhibition, as measured by displacement of 4 nM [3H]-dofetilide binding to HEK cells, and the predictions from a model of hERG pIC50 are shown for comparison.

If a predictive model does not yet exist for a property, it is first necessary to build an *in silico* model based on experimental data in order to use the Glowing Molecule. For example, the experimental results from molecules synthesised early in the project may be used. This modelling step captures the relationships between chemical structure and property value that is used to generate the Glowing Molecule. With the advent of robust, automatic modelling techniques it is no-longer necessary to have extensive computational chemistry experience to generate an *in silico* model for use in this way [11].

Conclusion

In this paper we have discussed how *in silico* and *in vitro* ADME data may be used, along with other relevant properties, to guide decisions in the selection and design of high quality molecules with a high likelihood of success.

The key to this process is to use all of the available data, experimental and predicted, to identify those molecules that achieve the best balance of properties for the project's objectives. We have presented a framework that enables the property requirements to be clearly defined and applied to prioritise molecules and chemical series. This explicitly takes into account the uncertainty in the underlying data to provide an objective assessment of the decisions that the data supports, enabling resources to be focused with confidence.

When ideal properties cannot be found in existing molecules, *in silico* models may be used to guide the design of chemical modifications to create a novel molecule with improved properties. The Glowing Molecule method described herein uses the information captured when building a model to highlight regions of a molecule that have a strong influence on the property, guiding medicinal chemists to focus efforts on the chemistry that is most likely to produce an improved molecule.

In order to get the most out of these technologies, it is essential that they be directly accessible to the key decision-makers, who are usually project scientists. This allows the impact of decisions or new ideas to be explored interactively; with the high pace of modern drug discovery, delays while waiting for a specialist to analyse potential options often means that a decision is taken before the results arrive. As most project scientists are not computational specialists, access to these algorithms must be via an intuitive interface to reduce the barrier to their uptake and maximise their effective use.

References

[1] J. A. DiMasi, R. W. Hansen, H. G. Grabowski, *J Health Econ.* **2003,** 22, 151
[2] H. Van de Waterbeemd, E. Gifford, *Nat. Rev. Drug Discov.* **2003**, 2, 192
[3] S. Ekins, B. Boulanger, P. W. Swaan, M. A. Hupcey, *J Comput Aided Mol Des.* **2002**, 16, 381
[4] M. D. Segall, A. P. Beresford, J. M. R. Gola, D. Hawksley, M. H. Tarbit, *Expert Opin. Drug Metab. Toxicol.* **2006**, 2, 325
[5] A. Krogh, *Nature Biotechnology* **2008**, 26, 195
[6] O. Obrezanova, G. Csanyi, J. M. Gola, M. D. Segall, *J. Chem. Inf. Model.* **2007**, 47, 1847
[7] A. Schwaighofer, T. Schroeter, S. Mika, J. Laub, A. T. Laak, D. Sulzle, U. Ganzer, N. Heinrich, K. R. Muller, *J. Chem. Inf. Model.* **2007**, 47, 407
[8] R. Burbidge, M. Trotter, B. Buxton, S. Holden, *Computers & Chemistry* **2001**, 26, 5
[9] M. Rowley, D. J. Hallett, S. Goodacre, C. Moyes, J. Crawforth, T. J. Sparey, S. Patel, R. Marwood, S. Patel, S. Thomas, L. Hitzel, D. O'Connor, N. Szeto, J. L. Castro, P. H. Hutson, A. M. MacLeod, *J. Med. Chem.* **2001**, 44, 1603
[10] M. Thomsen, S. Dobelb, P. Lassenc, L. Carlsend, B. B. Mogensene, P. E. Hansenf, *Chemosphere* **2002**, 49, 1317
[11] O. Obrezanova, J. M. Gola, E. J. Champness, M. D. Segall, J. Comput. Aided Mol. Des. **2008**, 22, 431