Addressing toxicity risk when designing and selecting compounds in early drug discovery

Matthew D. Segall[1] and Chris Barber[2]

[1] Optibrium Ltd, 7221 Cambridge Research Park, Beach Drive, Cambridge CB25 9TL, UK
[2] Lhasa Limited, 22-23 Blenheim Terrace, Woodhouse Lane, Leeds LS2 9HD, UK

Prioritising compounds with a lower chance of causing toxicity, early in the drug discovery process, would help to address the high attrition rate in pharmaceutical R&D. Expert knowledge-based prediction of toxicity can alert chemists if their proposed compounds are likely to have an increased likelihood of causing toxicity. We will discuss how multiparameter optimisation approaches can be used to balance the potential for toxicity with other properties required in a high-quality candidate drug, giving appropriate weight to the alert in the selection of compounds. Furthermore, we will describe how information about the region of a compound that triggers a toxicity alert can be interactively visualised to guide the modification of a compound to reduce the likelihood of toxicity.

Introduction

Toxicity of drugs and clinical candidates remains a significant issue for the pharmaceutical industry, leading to increased attrition and cost, late-stage failures and market withdrawals. Recent data from CMR-International [1] indicate that 22% of drug candidates entering clinical development in the period 2006–2010 failed owing to nonclinical toxicology or clinical safety issues. In preclinical development, toxicity and safety issues accounted for 54% of failures (18% of all preclinical candidates). These expensive late-stage failures account for a large proportion of the cost of pharmaceutical R&D, recently estimated to be US$1.8 billion per marketed drug [2].

For many marketed drugs, toxicity remains an issue, causing adverse drug reactions (ADRs) and leading to black-box warnings, restrictions on use and even withdrawals. These dramatically reduce or even eliminate the return on R&D and marketing investments and harm the reputation of pharmaceutical companies and the industry as a whole. A study by Lasser *et al.* [3] indicates that, of 548 new chemical entities approved by the FDA between 1975 and 2000, 10.2% acquired one or more black-box warnings and 2.9% were withdrawn. Recent, high-profile examples of market withdrawals include cerivastatin (2001), valdecoxib (2005, USA) and rosiglitazone (2010, Europe). Of particular concern are idiosyncratic ADRs that, owing to their rare occurrence, are unlikely to be detected during clinical trials.

From the sobering statistics above, it is clear that addressing failures due to toxicity would have a dramatic effect on the productivity of pharmaceutical R&D and the quality of the resulting drugs. Some toxicity is driven by the biological mechanism of the intended pharmacological action of a compound, particularly in the case of compounds intended for new targets for which the association with a therapeutic indication has not yet been validated. However, a significant proportion of observed toxicities are caused by unintended effects unrelated to the primary biological target. In the latter cases, it should be possible to reduce risk by focusing on structural motifs that are less likely to cause toxicity due to known mechanisms. Alternatively, if a likelihood of toxicity being observed in the clinic can be identified early in the process, *in vitro* or *in vivo* experiments can be prioritised to assess this risk before additional, downstream investments are made.

In the mid-1990s, a similar observation was made regarding a high rate of failure as a result of poor compound pharmacokinetics (PK) in clinical trials [4]. This led to the introduction of *in vitro* assays for high-throughput measurement of ADME properties in early drug discovery [5] and development of computational, or *in silico*, methods for the estimation of these properties [6,7]. The result has been a reduction in the proportion of clinical failures as a result of PK issues from an estimated 39% in 1991

Corresponding author:. Segall, M.D. (matt.segall@optibrium.com), (matt@optibrium.com)

1359-6446/06/$ - see front matter http://dx.doi.org/10.1016/j.drudis.2014.01.006

to approximately 10% in 2000 [8]. Unfortunately, during the same period, the overall failure rate was unchanged and the proportion of clinical failures attributed to toxicity or safety issues increased from approximately 14% to 30%. This, in turn, has motivated a recent trend in developing and introducing *in vitro* assays earlier in the drug discovery process, to identify potentially toxic compounds and halt their progression. Similarly, *in silico* methods for the prediction of toxicity can help to guide the design and selection of compounds with reduced risk of toxicity.

This article will focus on knowledge-based methods for prediction of toxicity (also described as rule-based) that produce a semi-quantitative estimate of toxicity hazards, based on experimental precedence for similar compounds. A number of expert systems have been developed that provide a rule-based approach to toxicity [9]. Other approaches, broadly described as statistical methods, rely on fitting a mathematical model of compound characteristics to empirical data using a variety of techniques including Support Vector Machines, Naive Bayes, Decision Trees and Random Forest [10–14]. The output from knowledge-based and statistical methods is the classification of compounds as toxic or otherwise or predictions of a numerical measure of toxicity (e.g. LD_{50}). The principles that we discuss herein for the application of *in silico* methods to address toxicity in early drug discovery can apply equally to both approaches.

In the following sections we will describe the principles of knowledge-based prediction of toxicity and the challenges posed by application in early drug discovery. We will discuss how these methods can be applied to the selection of compounds, giving appropriate weight to predictions of toxicity against other important factors, and provide feedback on strategies for redesign of drug candidates to reduce toxicity risk. Finally, we will present two applications of knowledge-based toxicity predictions – one for recently approved drugs and the other in the context of a hypothetical hit-to-lead project – before drawing some conclusions.

Knowledge-based prediction of compound toxicity

Expert knowledge-based predictive systems for small molecules are designed to emulate the decision-making process of a group of experts by applying a form of artificial intelligence whereby a knowledge base of facts is used to make a prediction by inferring relationships between facts through a process known as reasoning [15,16]. This enables the introduction of associated data such as reactivity or knowledge of the mechanism of action, and can cope with uncertainty and conflicting data that are common in the field of toxicity prediction. By contrast, purely statistical approaches derive probabilities of toxicity by taking a dataset of compounds, identifying descriptors that show a correlation to activity and use this to predict the toxicity of novel compounds. Statistical systems have the advantage of being fast to implement and can more efficiently cope with large datasets when the endpoint is relatively simple. Expert systems are particularly well suited to making predictions for toxicities derived through multiple mechanisms for which only incomplete datasets are available. Expert systems can often provide more interpretable predictions with detailed supporting documentation [9,17].

In silico systems in the field of toxicity typically predict hazard – the possibility of a chemical causing harm [18]. Expert systems frequently also provide an indication of the likelihood for the prediction to be correct, supporting evidence and a reasoned argument for the cause of the hazard, which might include an expert analysis, a mechanistic explanation or even an adverse outcome pathway (AOP) [19]. Although valuable, such predictions normally require further analysis to derive the risk – the probability of that toxicity being observed. A key part of that analysis is to determine the exposure of the chemical at the site of toxicity – a step that requires an understanding of the dosing regimen, the pharmacokinetics and potentially relevant biological details such as species, age, disease state, sex and the potential for drug–drug interactions. This means that a hazard prediction has to be considered in the context of a number of other factors to derive an assessment of risk.

BOX 1

Examples of the reasoning levels within Derek and their definitions

Certain	The proposition (prediction) is known to be true
Probable	There is at least one strong argument for the proposition and none against it
Plausible	The weight of evidence supports the proposition
Equivocal	There is an equal weight of evidence for and against the proposition

The Derek prediction engine (http://www.lhasalimited.org/) [20], applied in the examples below, provides a prediction (active/inactive) for each toxicity endpoint. If no evidence of toxicity has been found then 'No report' (nothing to report) is returned. A prediction of activity is typically associated with a structural alert, identifying the motif triggering the positive prediction, along with an associated likelihood. The likelihood qualifies this prediction; some of the likelihoods relating to positive predictions are shown in Box 1. In practise, it has been demonstrated that likelihood can be taken as a level of confidence because it correlates well with the accuracy of a prediction [21].

Expert systems are frequently applied in the later stages of drug development [22,23], where it might be necessary to produce an assessment of risk suitable for regulatory acceptance or to design *in vivo* studies that should be undertaken to support a submission. In such cases, features including mechanistic interpretation, expert commentary, documentation, performance statistics and supporting data are particularly valuable. At this stage of the process, relatively few compounds are assessed for toxicity and the endpoints can be relatively complex, meaning that training sets for *in silico* models tend to be sparse and do not always sufficiently capture the different mechanistic pathways at work. To overcome this, collaborative data sharing, through organisations such as Lhasa Limited, enables participating companies to gain knowledge of toxicities from proprietary data without revealing confidential information such as biological targets or chemical structures.

By contrast, these methods have been less commonly applied in early drug discovery, where the numbers of compounds considered are much larger and the scientists using the predictions are less likely to be expert toxicologists. This makes detailed examination of each prediction, using detailed supporting information,

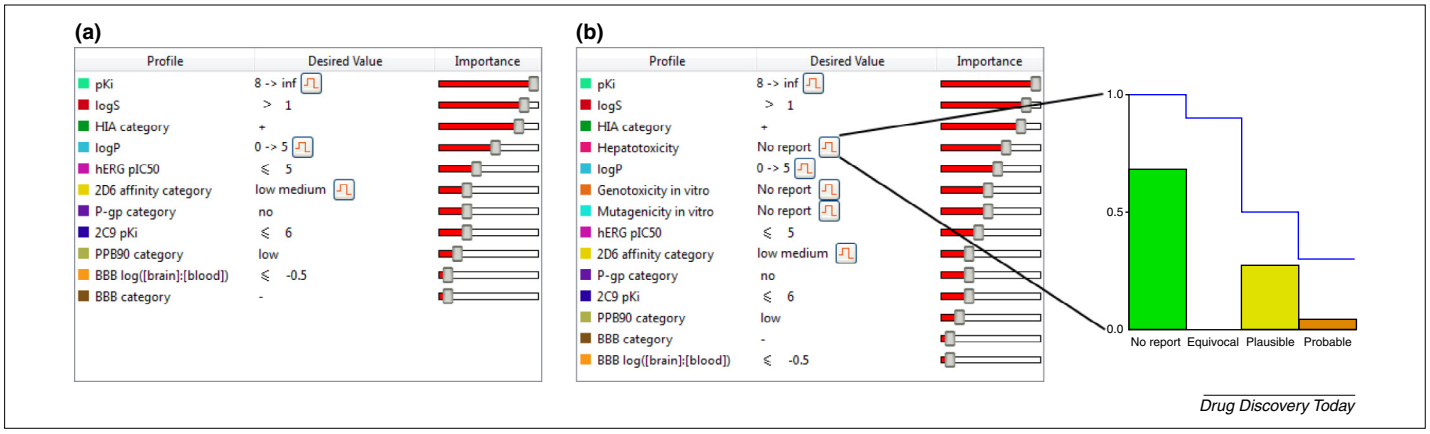

FIGURE 1

Example scoring profiles defining the ideal criteria (labelled 'desired value') for a range of experimental and predicted properties and the importance of each individual criterion to the overall objective of the project, specifically an orally dosed compound intended for a peripheral target. **(a)** An example of a profile that includes experimental potency against the target and predicted ADME properties. **(b)** Illustrates a profile combining these properties with knowledge-based predictions of toxicity endpoints. Also shown in (b) is an expansion of the criterion for hepatotoxicity, demonstrating how the impacts of different predicted likelihoods for this toxicity on the chance of a compound's success can be reflected by a 'desirability function' (blue line). On this graph, the desirability of each outcome is shown by the blue line and the scale on the y axis indicates the desirability on a scale of 0–1, where 1 indicates the ideal outcome. The histogram shows the distribution of the different predictions in the current dataset.

impractical. In this scenario, toxicity predictions must be appropriately integrated into decision-making processes to provide intuitive guidance on reducing toxicity risk and facilitate collaboration with expert toxicologists where expert guidance is required.

Guiding compound design and selection

Balancing toxicity with other factors

A high-quality drug must simultaneously satisfy many property requirements. Primary among these is achieving sufficient potency against the intended therapeutic target(s); however, to be safe and efficacious, a successful compound must also have appropriate ADME properties and, of course, avoid causing toxic effects at a therapeutic dose. Therefore, identifying high-quality lead and candidate compounds is a delicate balancing act, often described as multiparameter optimisation (MPO) [24].

Predictions of toxicity hazards must be balanced against other properties and given appropriate weight in the selection and design of compounds in early drug discovery. As discussed above, knowledge-based methods for toxicity prediction indicate if a compound has an increased likelihood of toxicity, but a toxicity alert is not a guarantee that a compound will be toxic. Therefore, one would give priority to compounds with no indications of toxicity over those with an alert, all other factors being equal; however, an alert might not be a sufficient reason to 'kill' a compound that meets many other requirements. The cost of incorrectly rejecting a good compound based only upon an uncertain prediction can be high, particularly in the absence of alternative options or if methods to mitigate the risk (such as a change to the dosing regimen) have not been considered.

Methods for MPO, such as probabilistic scoring [25], allow a project team to define a profile of property criteria that they require in an ideal compound. Furthermore, as illustrated in Fig. 1, each property criterion can be assigned an importance to reflect the impact of a property outcome on a compound's chance of success. The results of predictions or experimental property measurements for each compound are then assessed against the profile to generate a score representing the compound's likelihood of success (i.e. the probability of achieving an ideal property profile). This allows compounds with the best chance of downstream success to be effectively prioritised. Furthermore, the uncertainty in each property value, owing to experimental variability or statistical errors in predictions, can be explicitly taken into account to estimate the uncertainty in the overall scores. This, in turn, makes it clear when compounds can be confidently distinguished, based on the available data, and avoids inappropriate rejection of compounds based on an uncertain prediction or measurement.

Guiding compound redesign

An advantage of a knowledge-based approach to toxicity prediction is that the structural feature of a compound that is associated with an increased likelihood of toxicity is identified. This contrasts with many 'black box' statistical methods that provide a prediction with no feedback regarding the underlying relationship to the compound structure. Highlighting this alert on the structure of the compound provides valuable information for medicinal chemists considering optimisation strategies. Coupled with predictive models of other properties and an MPO method in an interactive environment, this information can be used to guide the design of an alternative compound to reduce the risk of toxicity without having a negative effect on other required properties. Fig. 2 shows an example of such an 'interactive designer'.

Analysis of recent drug approvals

To assess the potential value of applying knowledge-based toxicity prediction, all small molecule drugs approved in 2012 by the FDA Center for Drug Evaluation and Research (http://www.fda.gov/downloads/Drugs/DevelopmentApprovalProcess/HowDrugsareDevelopedandApproved/DrugandBiologicApprovalReports/UCM342733.pdf) were analysed against available endpoints in the Derek Nexus module for StarDrop (http://www.optibrium.com/).

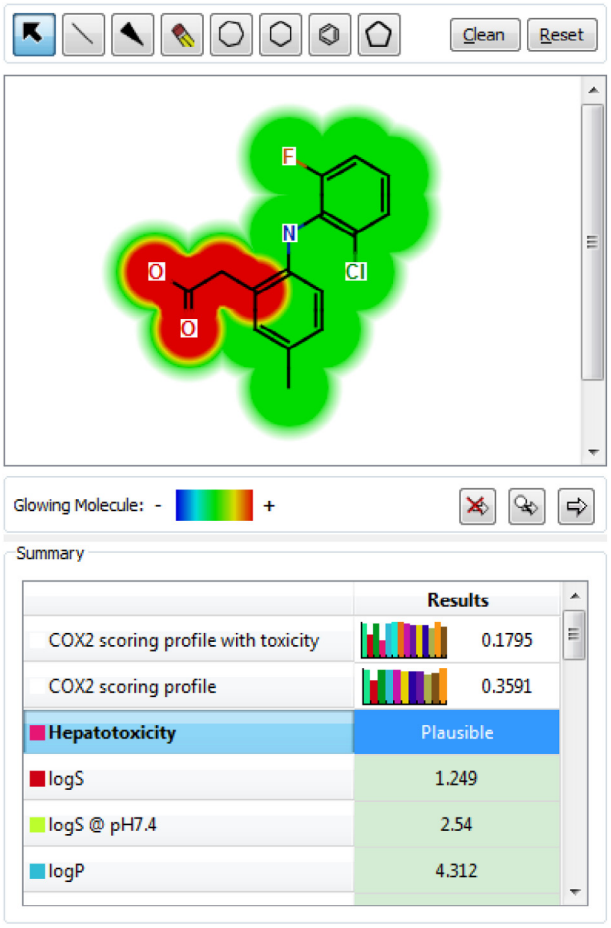

FIGURE 2

An example of an interactive designer in which the structural alert giving rise to the prediction of an increased chance of hepatotoxicity for lumiracoxib is highlighted in red. Such an environment enables exploration of strategies to reduce toxicity risk while providing instant feedback on the predicted impact of structural changes on multiple, relevant properties. *Abbreviation*: COX2, cyclooxygenase 2.

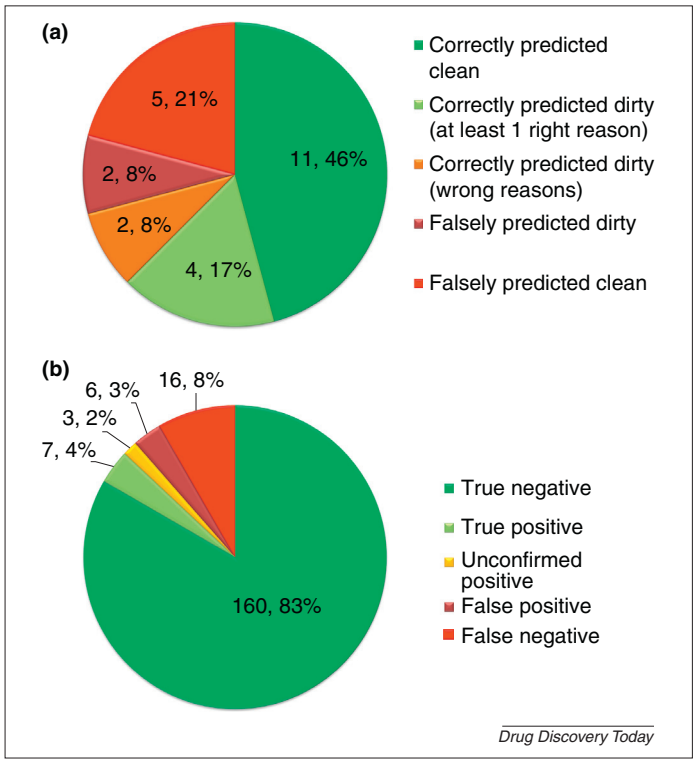

FIGURE 3

Results of predictions from the Derek Nexus module for StarDrop on the 24 compounds approved by the FDA in 2012. **(a)** The analysis on a per compound basis and **(b)** the analysis on a per endpoint basis.

The structures of these compounds were obtained from PubChem or NCI and regulatory label information was obtained from the FDA and/or European medicines Agency (EMA) to identify clinically observed toxicities together with relevant black-box warnings. Three endpoints: skin sensitisation and irritation of the eye or skin, were subsequently removed from the analysis because only one compound was topically administered hence predictions for these adverse events could not be validated. It should however be noted that the single topically administered compound, ingenol mebutate, was correctly predicted as a skin sensitiser. The predictions covered a range of important endpoints including hepatotoxicity, hERG-channel inhibition, developmental toxicity, teratogenicity, chromosomal damage (*in vitro* and *in vivo*), mutagenicity (*in vitro*) and carcinogenicity. All of the alerts returned were at the plausible level (meaning that the weight of evidence is for activity to be observed). This full dataset is available as Supplementary information. Of the limited set of 24 compounds, 11 were correctly predicted clean, six were correctly predicted with toxicities, five were falsely predicted clean and two were falsely predicted to have toxicities that were not observed. This is summarised in Fig. 3a. For the dataset of 24 compounds across eight endpoints a total of 16 predictions of toxicity were made, showing a sensitivity of 55% and a specificity of 85%. The breakdown of those alerts is shown in Fig. 3b.

Looking at the performance of individual alerts, three complex high-level endpoints were responsible for the majority of false predictions. The first, hepatotoxicity is a challenging endpoint to predict because there are many causes of liver toxicity, and many possible reasons why it might not be observed. One compound, aclidinium, was predicted to show hepatotoxicity but this is administered as a small inhaled dose making hepatotoxicity an unlikely event [26]. Two compounds (bosutinib and bedaquiline) were not predicted to show the observed hepatotoxicity but these were dosed at very high levels (>400 mg daily) and, in the latter case, toxicity was only shown when co-administered with other drugs that inhibit cytochrome P450 CYP3A4, the major clearance route for bedaquiline (http://www.accessdata.fda.gov/drugsatfda_docs/label/2013/204384s002lbl.pdf). Large doses of hepatically cleared compounds increase the risk of liver toxicity through the saturation of processes or the build-up of metabolites [27]. Developmental toxicity currently has only a limited number of alerts in the Derek engine and in this dataset only one out of the seven observed instances was correctly predicted. Three of the 'false negative' compounds were kinase inhibitors, which could indicate a lack of historical data from which to build good models because kinase inhibitors represent a relatively new class of drugs. There is however, growing evidence of a relationship between kinase inhibition and chromosomal and/or developmental toxicity [28], which should support further development of this alert. A third endpoint, teratogenicity, is incompletely understood, complex and driven by an array of

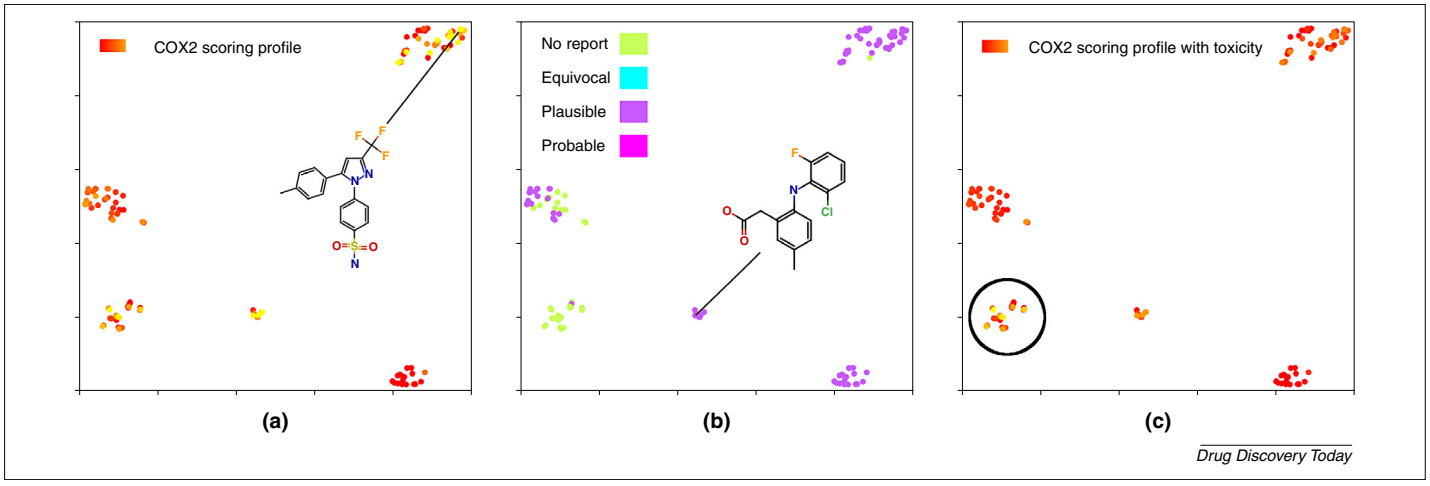

FIGURE 4

These 'chemical space' plots illustrate how predictions of the potential to cause toxicity can be combined with other experimental and predicted data to guide the selection of lead series in early drug discovery. Each point in a chemical space represents a single compound and the proximity of points indicates their structural similarity; 2D path-based similarity calculated by a Tanimoto index [30]. **(a)** The compounds in a library of compounds with cyclooxygenase 2 (COX2) inhibition data containing five clusters of similar compounds, coloured by compound score from red (low) to yellow (high). The score was calculated using the profile shown in Fig. 1a, taking into account only potency and ADME properties. From this it can be seen that multiple clusters contain compounds with high-scoring compounds. For reference, the point corresponding to celecoxib is identified. **(b)** The points coloured by predicted likelihood of hepatotoxicity, from which it can be seen that many regions of chemistry are predicted to have increased likelihood of exhibiting hepatotoxicity. The point corresponding to lumiracoxib, a known hepatotoxin, is highlighted in this plot. In **(c)** this information is combined with the data for compound potency, predicted ADME properties and predictions for mutagenicity and genotoxicity using the scoring profile shown in Fig. 1b. The colours indicate low-scoring compounds in red and high-scoring compounds in yellow, and the cluster containing the majority of high-scoring compounds is circled.

pathways – for example, lomitapide exhibited teratogenicity in rats and ferrets but not in rabbits. Teriflunomide is believed to be teratogenic through its primary mechanism of action – inhibition of dihydroorotate dehydrogenase – an essential enzyme for nucleotide synthesis (http://www.accessdata.fda.gov/drugsatfda_docs/label/2013/203858s002lbl.pdf).

This analysis suggests that knowledge-based toxicity predictions can be an effective tool to identify potential toxicities before a compound reaches the clinic. By flagging potential toxicities early in the drug development process, hazards can be assessed through early screening before significant investments have been made and by applying MPO analyses, these risks can be balanced against the potential benefits a drug might provide. This is highlighted in the case of carfilomib which despite giving six alerts (four of which were confirmed and two were not assessed; see Supplementary information), has been accepted as a treatment for cancer when other treatments are unsuccessful.

Application in early drug discovery

To illustrate one workflow for the practical application of these methods in the context of a hit-to-lead project, we have used a public domain dataset, derived from the ChEMBL database (https://www.ebi.ac.uk/chembl/). This dataset contains 152 compounds from multiple chemical series for which the inhibition of the cyclooxygenase 2 (COX2) enzyme has been determined experimentally, including the drugs celecoxib and lumiracoxib. This is typical of a dataset containing primary screening data in a hit-to-lead project targeting a fast-follower for an existing drug.

Fig. 4a shows the 'chemical space' of this library, in which the colour of a point represents the score of each compound against the scoring profile shown in Fig. 1a, including the experimentally measured target inhibition and a range of predicted ADME properties, but not considering predicted toxicity. This illustrates the distribution of the compound scores across the chemical diversity of the library and indicates that there are three clusters of similar compounds that are likely to yield compounds with a good balance of potency and ADME properties. These high-scoring compounds include the drugs celecoxib and lumiracoxib.

The potential for these compounds to cause toxicities was then predicted using the Derek Nexus module for StarDrop for endpoints including mutagenicity, hepatotoxicity and genotoxicity. Mutagens cause heritable changes to DNA, whereas genotoxins damage a cell's genetic material but do not necessarily cause permanent damage to DNA sequences. Fig. 4b shows the prediction of hepatotoxicity mapped onto the chemical space of the COX2 library, which clearly shows that several of the clusters have plausible evidence of hepatotoxicity and should be considered with care. Among those compounds with evidence of hepatotoxicity is lumiracoxib, which was withdrawn from the market in several countries, mostly as a result of hepatotoxicity concerns, and has never been approved for use in the USA.

The toxicity predictions can be combined with the *in vitro* and *in silico* data for other properties in an overall scoring profile (Fig. 1b) giving appropriate weight to the predictions of toxicity against the other factors. The resulting scores are plotted in the chemical space shown in Fig. 4c, in which one cluster clearly stands out as having several compounds with the highest likelihood of yielding a high-quality lead series with good ADME properties and reduced chance of toxicity.

It is noteworthy that celecoxib (the gold-standard COX2 inhibitor) [29] is also identified as having plausible evidence of toxicity, illustrating the importance of balancing the potential for toxicity against the benefits. One advantage of avoiding hard filters, by using a weighted scoring profile and taking into account

the uncertainty in the underlying data, is that the series including celecoxib and lumiracoxib would not be rejected outright. For example, the score for celecoxib (0.15 ± 0.08) is statistically not significantly different from the top-scoring compound (0.45 ± 0.30). This indicates that a rigorous strategy should select a small number of compounds from this series to confirm experimentally the required properties before making a final choice of lead series.

Finally, considering the structure of lumiracoxib in Fig. 2, a single functionality is highlighted as the cause of the structural alert for increased hazard of hepatotoxicity, in common with all other members of this series. This suggests that approaches for reducing the associated risk, while retaining potency and other desirable properties, can be investigated at an early stage before rejecting this class of compounds.

Concluding remarks

A key strategy to reduce the long timelines and spiralling cost of pharmaceutical R&D is to target safe and efficacious compounds as early as possible in the drug discovery process. Taking all available information into account, from predictive and experimental sources, as early as possible, increases the likelihood of delivering a high-quality lead and, ultimately, a development candidate with an improved chance of success in the clinic. Furthermore, a lead series with a good balance of properties is less likely to require many, long and costly lead optimisation cycles; an important factor identified to reduce the overall cost per marketed drug [2].

Knowledge-based prediction of toxicity has an important part to play in this process, guiding the selection and optimisation of compounds when *in vitro* and *in vivo* toxicity data are often not available, owing to the high cost and long timescales of experimental measurements. However, as with any predictive method, the uncertainties in the predicted outcomes should be taken into account and appropriate weight should be given to these results, relative to other property requirements for a high-quality compound for a drug discovery project's objective. In this article, we have illustrated how toxicity predictions can be incorporated into an MPO approach to identify compounds quickly that have an appropriate balance of properties and guide the optimisation of compounds with potential liabilities.

Acknowledgements

The authors would like to thank Ed Champness and Nick Foster for useful discussions and comments on the draft manuscript. They would also like to thank the software teams at Optibrium and Lhasa Limited for the development of the methods described herein, in particular Chris Leeding and Paul Cooke for developing the Derek Nexus module in StarDrop. Derek KnowledgeBase 2012 1.0 was used for this analysis.

Appendix A. Supplementary data

Supplementary material related to this article can be found, in the online version, at http://dx.doi.org/10.1016/j.drudis.2014.01.006.

References

1 CMR International Global R&D Performance Metrics programme, CMR International, a Thomson Reuters business.
2 Paul, S.M. *et al.* (2010) How to improve R&D productivity: the pharmaceutical industry's grand challenge. *Nat. Rev. Drug Discov.* 9, 203–214
3 Lasser, K.E. *et al.* (2002) Timing of new black box warnings and withdrawals for prescription medications. *JAMA* 287, 2215–2220
4 Kennedy, T. (1997) Managing the drug discovery/development interface. *Drug Discov. Today* 2, 436–444
5 Tarbit, M.H. and Berman, J. (1998) High-throughput approaches for evaluating absorption, distribution, metabolism and excretion properties of lead compounds. *Curr. Opin. Chem. Biol.* 2, 411–416
6 Van de Waterbeemd, H. and Gifford, E. (2003) ADMET *in silico* modelling: towards prediction paradise? *Nat. Rev. Drug Discov.* 2, 192–204
7 Ekins, S. *et al.* (2001) Towards a new age of virtual ADME/TOX and multidimensional drug discovery. *J. Comput. Aided Mol. Des.* 16, 381–401
8 Kola, I. and Landis, J. (2004) Can the pharmaceutical industry reduce attrition rates. *Nat. Rev. Drug Discov.* 3, 711–716
9 Dearden, J.C. *et al.* (1997) The development and validation of expert systems for predicting toxicity: the report and recommendations of an ECVAM/ECB workshop (ECVAM Workshop 24). *ATLA* 25, 223–252
10 Roncaglioni, A. *et al.* (2013) In silico methods to predict drug toxicity. *Curr. Opin. Pharmacol.* 13, 802–806
11 Arodz, T. *et al.* (2006) Ensemble of linear models for predicting drug properties. *J. Chem. Inf. Model.* 46, 416–423
12 Nigsch, F. and Mitchell, J.B.O. (2008) How to winnow actives from inactives: introducing molecular orthogonal sparse bigrams (MOSBs) and multiclass Winnow. *J. Chem. Inf. Model.* 48, 306–318
13 Nigsch, F. *et al.* (2008) Ligand–target prediction using Winnow and naive Bayesian algorithms and the implications of overall performance statistics. *J. Chem. Inf. Model.* 48, 2313–2325
14 Tropsha, A. (2012) Recent trends in statistical QSAR modeling of environmental chemical toxicity. In *Molecular, Clinical and Environmental Toxicology*, (Vol. 3) (Luch, A., ed.), Springer
15 Greene, N. *et al.* (1999) Knowledge-based expert systems for toxicity and metabolism prediction: DEREK, StAR and METEOR. *SAR QSAR Environ. Res.* 10, 299–314
16 Dearden, J.C. (2003) In silico prediction of drug toxicity. *J. Comput. Aided Mol. Des.* 17, 119–127
17 Cronin, M. and Madden, J., eds) (2010) *In Silico Toxicology: Principles and Application*, RSC Publishing
18 Greene, N. (2002) Computer systems for the prediction of toxicity: an update. *Adv. Drug Deliv. Rev.* 54, 417–431
19 Ankley, G.T. *et al.* (2010) Adverse outcome pathways: a conceptual framework to support ecotoxicology research and risk assessment. *Environ. Toxicol. Chem.* 29, 730–741
20 Marchant, C.A. *et al.* (2008) In silico tools for sharing data and knowledge on toxicity and metabolism: Derek for windows, Meteor, and Vitic. *Toxicol. Mech. Methods* 18, 177–187
21 Judson, P.N. *et al.* (2013) Assessing confidence in predictions made by knowledge-based systems. *Toxicol. Res.* 2, 70–79
22 Sutter, A. *et al.* (2013) Use of *in silico* systems and expert knowledge for structure-based assessment of potentially mutagenic impurities. *Regul. Toxicol. Pharmacol.* 67, 39–52
23 Dobo, K.L. *et al.* (2012) In silico methods combined with expert knowledge rule out mutagenic potential of pharmaceutical impurities: an industry survey. *Regul. Toxicol. Pharmacol.* 62, 449–455
24 Segall, M.D. (2012) Multi-parameter optimization: identifying high quality compounds with a balance of properties. *Curr. Pharm. Des.* 18, 1292–1310
25 Segall, M.D. *et al.* (2006) Focus on success: using *in silico* optimisation to achieve an optimal balance of properties. *Expert Opin. Drug Metab. Toxicol.* 2, 325–337
26 Lasseter, K. *et al.* (2012) Safety and pharmacokinetics of multiple doses of aclidinium bromide administered twice daily in healthy volunteers. *Pulm. Pharmacol. Ther.* 25, 192–199
27 Wang, K. *et al.* (2009) Circulating microRNAs, potential biomarkers for drug-induced liver injury. *Proc. Natl. Acad. Sci. U. S. A.* 106, 4402–4407
28 Olaharski, A.J. *et al.* (2010) Identification of a kinase profile that predicts chromosome damage induced by small molecule kinase inhibitors. *Toxicol. Sci.* 118, 266–275
29 Moore, R.A. *et al.* (2005) Tolerability and adverse events in clinical trials of celecoxib in osteoarthritis and rheumatoid arthritis: systematic review and meta-analysis of information from company clinical trial reports. *Arthritis Res. Ther.* 7, R644–R665
30 Rogers, D.J. and Tanimoto, T.T. (1960) A computer program for classifying plants. *Science* 132, 1115–1118