

In silico ADME-Tox prediction: The more, the merrier

Gregory M Banik

Bio-Rad Laboratories, Philadelphia, PA, USA



In silico, or virtual, ADME-Tox prediction is a method that can be used to evaluate the ADME-Tox profile of a compound, even before it is synthesized, and thus to concentrate finite resources and energy on those few compounds most likely to succeed. The result can dramatically lower costs and reduce cycle time for pharmaceutical R&D.

Problems with a compound's absorption, distribution, metabolism, excretion, or toxicity (ADME-Tox) have been identified as a principal cause of failure in late-stage pharmaceutical R&D. The numerous international conferences and market research reports in 2003 and 2004 focused on ADME-Tox underscore the continuing importance of this issue. Sinead Igoe, Frost & Sullivan's Industry Analyst for Drug Discovery Technologies and author of the recent *World ADME/Tox in Drug Development Markets* report, estimates that the worldwide spending on in silico ADME-Tox technologies in 2004 will be \$46 million, with high double-digit growth over the next several years. This nascent market has attracted a number of companies providing in silico ADME-Tox prediction tools, and while the focus and attention these tools is both timely and warranted, much remains to be done in terms of integrating the existing tools into a single, consistent workflow environment.

Multiple endpoints

Ekins et al outlined a set of computational ADME-Tox endpoints required in drug discovery (Figure 1) as well as a set of requirements for the integration of these endpoints with in vitro data (Figure 2). Together, they form an overall statement of requirements for an in silico ADME-Tox prediction system.

The use of models to cover multiple ADME-Tox endpoints will increase the overall breadth of coverage. An informatics

Primary models

- Solubility
- Absorption
- Mutagenicity
- Bioavailability
- Metabolic stability
- Blood-brain barrier permeability
- Cardiac toxicity (hERG)
- Plasma protein binding

Secondary models

- Transport (uptake and efflux)
- General toxicity
- Hepatotoxicity
- Nephrotoxicity
- Immunogenicity
- Neurotoxicity
- Drug-drug interactions (cytochrome P450)

Figure 1. Computational ADME-Tox models required for drug discovery.

system mapped to the in silico ADME-Tox workflow process to cover multiple endpoints simultaneously would need to include tools for validation, prediction, analysis, design, and reporting (Figure 3).

• First and foremost, the models must be validated. Without statistically rigorous validation, the models will always be suspect and therefore useless. Toward that end, an integrated application is required to validate any of the in silico ADME-Tox models with a set of structures and experimental results. Validation is typically performed by an expert user such as a computational chemist or molecular modeler - someone familiar with good statistical practices.

• Once the models have been validated, a compound library (real or virtual) without measured ADME-Tox endpoints can be processed in batch fashion to create a library with predicted endpoints. The good news with in silico ADME-Tox prediction is that it is possible to generate large numbers of data points for large numbers of compounds. The bad news, of course, is that *it is possible to generate large numbers of data points for large numbers of compounds.*

1. A consistent source of in vitro data for each system modeled
2. A large, diverse data set of molecular structures tested with a range of activities
3. Multiple integrated molecular descriptor generation tools
4. Methods for removing correlating descriptors and missing values
5. Multiple algorithms
6. Statistical methods for testing, validating and selecting models for use
7. Providing the model/s in an accessible intuitive output
8. Integration of models with other available tools
9. Ensure the models are used at the appropriate point in drug discovery
10. Continual update and refinement of models with new in vitro data.

Figure 2. Computational ADME-Tox integration requirements.



Figure 3. In silico ADME-Tox workflow.

To turn this flood of unintelligible data into actionable information, one must resort to analysis and datamining.

- Datamining is defined as ‘an information extraction activity whose goal is to discover hidden facts contained in databases.’ ADME-Tox tools require project teams to assign the relative importance, not only of each endpoint, but also of the ranges of values for each endpoint. Results can be visualized for analysis in a color-coded representation often referred to as a ‘heatmap’, allowing an entire library of compounds to be rank-ordered using any combination of experimental or predicted values. This ordering process allows those compounds with the highest probability of success to be tested in slower and more expensive in vitro or in vivo tests first and those with the lowest probability of success to be tested last.
- The in silico ADME-Tox workflow is a cyclical rather than a linear process, and continues in an iterative fashion until a desirable outcome is achieved. The iterations come through re-design of the structures in the library in an attempt to maximize the potency of the molecule while minimizing any deleterious effects. In such a process, medicinal chemists desire a tool that displays the complete in silico ADME-Tox profile of a molecular structure and yet allows the structure to be modified quickly to display the complete in silico ADME-Tox profile of the modified structure. This type of ‘what if’

simulation allows for real-time assessment by the medicinal chemist of any proposed changes to their lead compounds.

- Throughout the workflow, the ability to generate quick yet informative reports is important. Reporting allows information to be recorded, exchanged, and communicated with peers and managers, increasing the efficiency of the workflow process and thereby reducing the overall cycle time.

Multiple models

With in silico ADME-Tox prediction, it is possible that a single model will provide acceptable results. Sometimes, however, a single model will not work. By definition, all models, including ADME-Tox models, are simulations of reality, and as such will never be completely accurate. When multiple models are combined as a single consensus model, however, more accurate predictions can be achieved. Recent research has demonstrated the benefits of employing multiple complementary models for the same ADME-Tox endpoint in a consensus modeling approach to provide significantly greater accuracy over that of a single model.

In general, three types of consensus models exist:

- Real variable: A real or continuous variable consensus model involves the combination of multiple models whose output is a real or continuous variable. Examples of real variable models include

log P, log D, plasma protein binding, or any other model whose output is a real number.

- Boolean variable: A Boolean or classification variable consensus model involves the combination of multiple models whose output is a Boolean or classification variable. An example of a Boolean variable model would be a mutagenicity model whose output is either ‘true’ (the molecule is predicted to be mutagenic) or ‘false’ (the molecule is predicted to be non-mutagenic).

- Mixed variable: If one model for an endpoint outputs a real variable while a different model for the same endpoint outputs a Boolean variable, the results cannot be combined without normalization to the same scale. This typically involves mapping one range of the real number model to ‘true’ while mapping a non-overlapping range for the same model to ‘false’. Once on the same Boolean scale, the models can be combined in a consensus fashion as with Boolean variables.

A recognized method for creating a real-variable consensus model starts with a set of k models and a set of z structures, each of which has an experimentally measured value related to the model of interest. A different in silico ADME-Tox prediction is performed by each model on each structure, resulting in a k by z matrix of predicted values, that is, k different predicted values for each structure. The k by z matrix of predicted values is compared to the experimental values, and an integrated N -fold cross-validation process is used to derive the consensus model.

In N -fold cross-validation, the set of predicted and actual values is split into N different groups. The first group is used as a test set, and the remaining N minus one sets are used as a training set. Using this test and training set, a set of weights is determined for each model to give the best correlation between the sum of the weighted predictions and the experimental values. This process is repeated using the second group as the test set and the remainder as the training set, and the process is repeated to give N sets of weights for each of the models. As a result of this process, each compound is in a test

set once and in a training set N minus one times. The N weights for each model are averaged, and it is these weights that are used in the weighted average of the resulting consensus model, resulting in a more predictive and cross-validated model.

Once created and cross-validated, the weights can be saved as consensus model and used in a high-throughput fashion to predict ADME-Tox values for structures without experimentally-measured ADME-Tox values, thus facilitating the *in silico* lead optimization process in pharmaceutical R&D.

For Boolean, or classification models, there are several types of consensus models that are available:

- In the worst-case scenario model, the consensus output will be 'false' if and only if all models in the consensus have a value of 'false'.

- The best-case scenario model is the converse of the worst-case

scenario model, outputting a value of 'true' 'if and only if all models in the consensus have a value of 'true'.

- The majority rules model returns a 'true' value if most of the models return a 'true' value and a 'false' value if most of the models return a 'false' value.

- The percent agreement model is similar to the majority rules model except that a threshold for a required percent agreement can be specified.

An example of a mixed mode consensus model mapped to a Boolean scale is that of mutagenicity predicted by CompuDrug's HazardExpert mutagenicity model and ChemSilico's CSGenotox mutagenicity model. The former model outputs a real variable probability of the molecule's mutagenicity while the latter outputs a '0' if the compound is predicted to be non-mutagenic and a '1' if the compound is predicted to be mutagenic. Mapping the output of the HazardExpert model greater than or equal to zero, but less than 50 to 'false' and mapping the output greater than or equal to 50, but less than or equal to 100 to 'true' allows the output of this model to be combined with that of the CSGenotox model. The results of a validation study for these two models using a

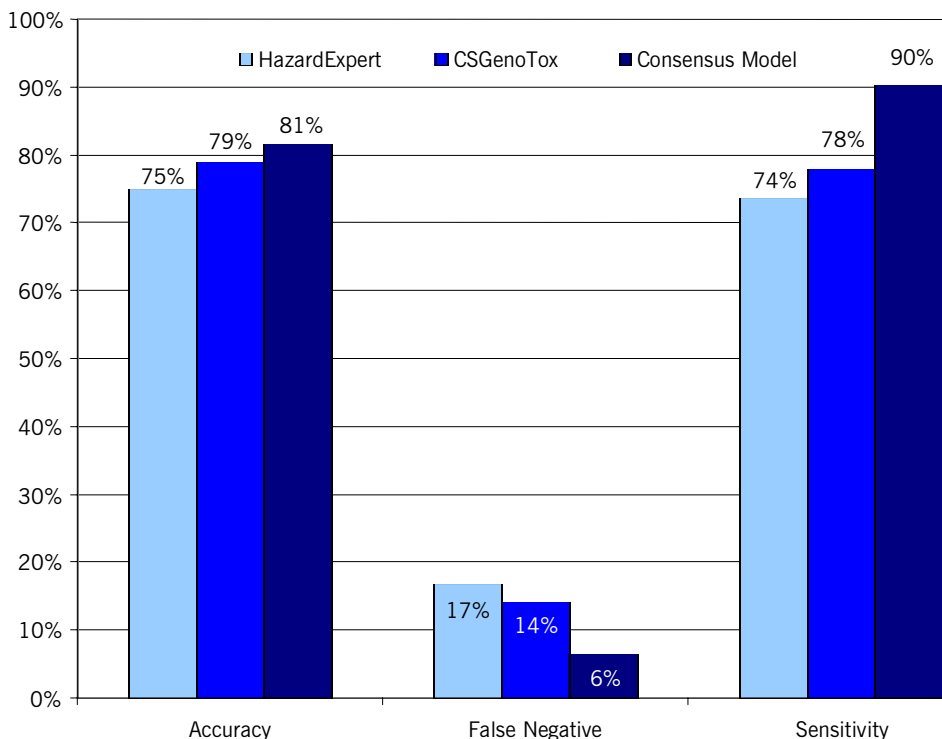


Figure 4. Classification consensus model results.

worst-case scenario consensus approach shows that the consensus model improves the overall accuracy of the model, the number of false negatives, and the sensitivity of the model, which is a measure of a model's ability to correctly identify high-risk compounds (Figure 4).

Finally, it is possible to create substructure-localized consensus models, that is, consensus models that have been fine-tuned according to the specific requirements of substructural classes of compounds. ADME-Tox models may predict some classes of compounds better than others. By taking the strengths of each model into account for each substructural class, a more accurate consensus model can be achieved for each class of interest.

Global models versus local models

While having access to global, commercially available models provides an 'out-of-the-box' solution for customers, there may be instances where either a global model or a localized consensus of global models will not perform as well as a true local model. This would certainly be the case for a model of biological activity that is trained on a company's own proprietary chemical

structures and biological activity data. Toward this end, the ability to add a company's own models to a complete ADME-Tox environment is key.

Partnering: The more, the merrier

Bio-Rad entered the *in silico* ADME-Tox arena in 2003 with the launch of the KnowItAll ADME-Tox Edition. Rather than attempt to develop the large number of predictive ADME-Tox models required, Bio-Rad chose to adopt a strategy of partnership and integration. Toward that end, our existing collaborations with ChemSilico, CompuDrug/ComGenex, Strand Genomics, and Syracuse Research have allowed us to offer the world's largest collection of ADME-Tox predictors in a single, unified software application.

The integration of our partners' technology into KnowItAll allows each predictor to take advantage of its unique informatics environment, including fully integrated tools for model validation, library profiling, data analysis, iterative compound redesign, and report generation. Depending on the relationship with the partner, the resulting product may be distributed exclusively by Bio-Rad, distributed exclusively by the partner, or by both.

Of particular importance to all our partners is the fact that there also exists a version of the KnowItAll software that is fully localized into Japanese.

The KnowItAll Software Developer's Kit allows pharmaceutical companies to incorporate their own internally-developed technologies into the same environment, bridging the global versus local model gap. In addition, combining all the informatics tools and ADME-Tox models in a single, intuitive interface dramatically diminishes a company's learning curve issues associated with utilizing multiple software packages and greatly reduces 'application proliferation'. In addition to the obvious benefit of access to multiple, complementary models, further gains can be realized by refocusing IT efforts away from homegrown ADME-Tox prediction systems, which are difficult to create and burdensome to maintain. Overall, there are significant benefits to be had by scientific and IT staff members in pharmaceutical companies who are being asked to do more with less.

Common sense dictates that relying on a panel of experts will provide more reliable results than relying upon a single voice. An initiative such as Bio-Rad's

ADME-Tox partnership program enables the combination and consolidation of proven technology in an informatics environment that can improve workflow and enhance productivity.

Summary

Recently, ADME-Tox has gained prominence as a significant bottleneck in pharmaceutical R&D. To address this issue, in silico ADME-Tox prediction offers a means to significantly reduce attrition in the R&D process and thereby increase the return on investment in R&D. New computational methods, including consensus modeling, show promise for increasing the accuracy of in silico ADME-Tox prediction used for virtual screening in pharmaceutical lead optimization. New business alliances, such as those being initiated by Bio-Rad, show promise for creating a consistent, unified approach for in silico ADME-Tox prediction that will deliver on the promise of 'the more, the merrier'.

Acknowledgements

I would like to thank my associates involved with our ADME-Tox software R&D, including Ty Abshear and Karl

Nedwed, and especially Michelle D'Souza and Yann Bidault for their considerable and elegant research efforts in consensus modeling and model normalization. I also offer my most sincere and deepest thanks to our collaborators, including Ferenc Darvas and László Üрге at ComGenex, Joseph Votano and Marc Parham at ChemSilico, Kalyanasundaram Subramanian at Strand Genomics, and Philip Howard, Bill Meylan, and Jay Tunkel at Syracuse Research, for their great assistance in this effort. Finally, I would like to thank Sinead Igoe from Frost & Sullivan for her kind assistance.

Gregory M Banik

General Manager
Bio-Rad Laboratories
Informatics Division
3316 Spring Garden Street
Philadelphia
PA 19104
USA

Email: gregory_banik@bio-rad.com
www.knowitall.com/adme

FURTHER READING

Igoe S (2003) **World ADME/Tox in drug development markets.** *Frost & Sullivan.*

Ekins S *et al* (2002) **Towards a new age of virtual ADME/TOX and multidimensional drug discovery.** *Journal of Computer-Aided Molecular Design* **16**:381-401.

Shimada J *et al* (2002) **Integrating computer-based *de novo* drug design and multidimensional filtering for desirable drugs.** *TARGETS* **1**:196-205.

Baurin N *et al* (2004) **2D QSAR consensus prediction for high-throughput virtual screening. An application to COX-2 inhibition modeling and screening of the NCI database.** *Journal of Chemical Information and Computer Sciences* **44**:276-285.

Banik GM (2003) **Integrating ADME/Tox prediction and informatics.** *Cambridge Healthtech Institute's Predictive ADME and Predictive Toxicology Conference*, Boston, MA, USA.

Mattioni BE, Kauffman GW, Jurs, PC (2003) **Predicting the genotoxicity of secondary and aromatic amines using data subsetting to generate a model ensemble.** *Journal of Chemical Information and Computer Sciences* **43**:949-963.

White A (2003) **Predictive ADME and toxicity modeling - An emerging role in high throughput screening and drug discovery.** *The Center for Business Intelligence Predictive ADME/Tox Conference*, Philadelphia, PA, USA.

Pearl GM, Livingston-Carr S, Durham SK (2001) **Integration of computational analysis as a sentinel tool in toxicological assessments.** *Current Topics in Medicinal Chemistry* **1**:247-255.

Kennedy T (1997) **Managing the drug discovery/development interface.** *Drug Discovery Today* **2**:436-444.