PEER REVIEWED

Guiding Drug Optimisation Using Deep Learning Imputation and Compound Generation

The use of machine learning (ML) methods is now commonplace in many disciplines and artificial intelligence (AI) is on the rise, promising better and smarter solutions to 'all your problems'. However, despite the hype, there is increasing evidence we have entered the next 'AI winter' or the so-called 'trough of disillusionment' in the ongoing hype cycle[1]. There is still a gap in understanding on the route from traditional and well-understood statistical modelling methods to the poorly-defined promises of AI, and exactly how *the majority* of researchers can cross that gap is not clear.

Researchers in drug discovery are familiar with quantitative structure activity relationship (QSAR) model building methods. Many of these methods now employ forms of machine learning (ML), a sophisticated form of 'fitting functions to data'. The question is how to leap forward from this well-known and comfortable ML space toward sophisticated AI tools, by which we mean: A connected set of ML components in an automated system which together produce a rich behaviour capable of solving complex tasks. The lesser known 'AI' is *augmented* intelligence, and there is no reason why a human cannot be part of the connected components in this sophisticated AI system. The combination of a human expert and superior tools has been found to be optimal as well as convenient[2].

We will describe the outcomes and discoveries made by connecting: A state-of-the-art data imputation method[3], in this case using deep learning[4,5]; generative methods based on machine learning[6] and evolutionary[7] algorithms; optimisation processes for goal-seeking; and probabilistic scoring[8], a form of multi-parameter optimisation (MPO) [9] for guiding compound prioritisation decisions, without resorting to harsh filtering methods[10]. We illustrate this with an example application finding a confidently active compound against a novel malaria target[11], and outline our future vision for these methods.

Method

The schematic for the 'AI' process used in this work is shown in Figure 1. Each block represents a component or process, colour-coded as either an input (data, or parameters), ML method, automated step (such as a script or advanced software process), or outputs, which amount to *confident predictions of active compounds with an appropriate property profile and confident identification of missed opportunities*. In general, humans will only have to interact with the inputs and outputs of the system, but this feedback is essential to get the best results out of the system, through augmentation of the AI with human expertise and vice versa.

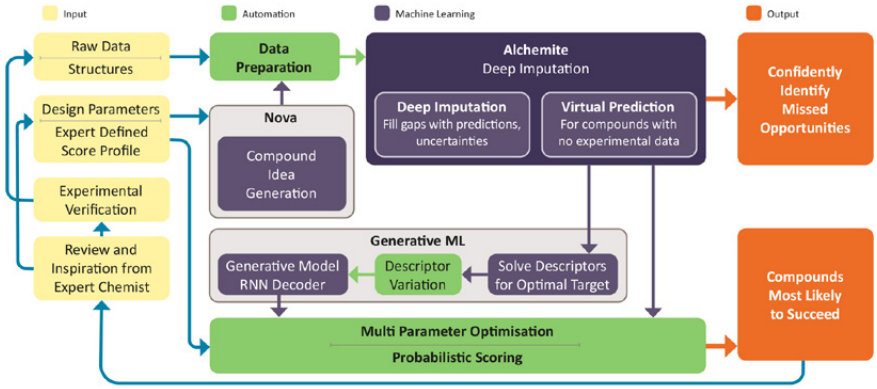

Figure 1: Schematic of the augmented/AI process used to generate confident predictions for virtual compounds. Components are colour-coded as inputs (yellow), machine learning tools (purple), automated steps (green) and outputs (orange).

For most applications, the starting point is sparse and noisy raw experimental data for existing compounds and their chemical structures (Figure 1, top left) [3]. For a typical drug discovery project, these data could be a combination of experimental assays for activities and absorption, distribution, metabolism and elimination (ADME) endpoints[5].

Data Preparation

First, we will follow the input data rightwards; the raw data are cleaned with automated routines and transformed into units that are more suitable for machine learning; for example, it is common to transform IC_{50} quantities to pIC_{50} [5] (the negative log of the IC_{50} in Molar units). Molecular descriptors are also generated for the compound structures as input; for this work, descriptors were generated, including whole-molecule properties such as logP, MW, TPSA and SMARTS based fragment matches, but in principle any descriptors can be used at this step.

Modelling

The next step is critical: the sparse and noisy data are *imputed* using a state-of-the-art deep learning method called Alchemite[4], which has seen great success in heterogeneous datasets from real drug discovery projects[5]. Modern imputation methods offer clear benefits over a standard predictive approach[3], they make better use of existing data, can handle sparse and noisy experimental results and provide robust uncertainty estimates for each missing value which is predicted[3,5]. From this point, in Figure 1 we can immediately find high-potential compounds that are in the training set, but only have partially measured experimental data. We can also "confidently identify missed opportunities" as illustrated as an output in the schematic, which is a valuable positive outcome[10]. Examples of this could be incorrect or inconsistent experimental data in the inputs, which the deep imputation model can highlight that they lie outside of the expected range based on the error bars in the prediction, enabling them to be flagged for retesting. When training this deep imputation model, we can also build a *virtual model* [3]. This can make predictions for virtual

compounds, i.e. those that have not yet been synthesised, with greater accuracy than the best QSAR methods[5], while simultaneously providing robust error bars. This virtual model forms the explorative basis for the AI allowing it to judge new compounds which are generated by two methods.

Compound Generation

Broadly speaking, there are two kinds of approach for compound generation, a *bottom up* and a *top down* approach:

Bottom up: Start from known chemistry and optimise outward to explore related compounds to search for those that are likely to satisfy project requirements (activity, selectivity, ADME, etc.). This method closely resembles a human-led drug discovery programme, starting from known hits and performing synthetically accessible variations around promising compounds, but can explore many more, diverse ideas in a shorter time than even an expert chemist[12]. In algorithmic terms, it can be slow to iteratively approach the goals, but the proposed chemistry will be familiar to project chemists and more likely to be synthesisable. In Figure 1, this is covered by the component 'Nova™', a module in StarDrop™, which uses advanced evolutionary algorithms to generate libraries of virtual compounds using realistic chemistry transformations from the medicinal chemistry literature[7].

Top down: This strategy attempts to generate descriptors for an 'ideal compound' which would be predicted to have the desired properties. The algorithm then generates a structure that matches these ideal descriptors. This method is an example of generative ML methods[6] and is shown as '*Generative ML*' in Figure 1. If the model predictions are accurate and a solution can be found, then the compound is likely to fulfil all of the project requirements. However, these methods often struggle to make *synthetically accessible* and *drug-like* compounds, and models without uncertainty estimates may give untrustworthy answers. The deep imputation method used in this work does provide uncertainty estimates, and these can be factored into the optimisation process. For our implementation, we solve for the ideal descriptor vector using a gradient descent optimisation layer over the Alchemite model. This layer varies the *descriptor* while minimising the difference in the predictions of the fixed model and the desired properties for a compound. The solved descriptors are cleaned and minor variations are made about the solution. These idea descriptor vectors are subsequently entered into a recurrent neural network (RNN) decoder[6] which is trained to write out SMILES representations as compound suggestions which meet the input descriptor profile. In our case, the StarDrop descriptors were generated for the original dataset alongside a large portion of the ChEMBL database. This meant the RNN would generate SMILES similar to ChEMBL compounds which also match the target ideal descriptors used as input.

Probabilistic Scoring

An important step is to take all of the virtual compounds generated from both methods, along with their predictions and uncertainties, and apply MPO to prioritise them for further consideration. This is because a high-quality compound must exhibit not only activity but also an appropriate balance of physicochemical and ADME properties. In this AI application we use the probabilistic scoring method[8] in which an experienced user can define a profile of property criteria that represents the desired outcomes for an ideal compound. The algorithm estimates the likelihood of success of each compound, taking the uncertainties in the property values into account. This enables the generated compounds to be prioritised and the highest scoring shown to a human expert. The most promising can be taken forward for synthesis and experimental studies, or the expert can update the scoring profile or the design parameters for the Nova module to generate a new list of suggestions.

Application

A practical application of the AI process shown in Figure 1 was demonstrated for the Open Source Malaria (OSM) challenge, where various teams were invited to build predictive models for pfATP4 activity based on an open source dataset[11]. The data were sparse and noisy and compounds had been measured across different labs, protocols and sensitivities, and assays had been performed on different strains of drug-resistant *Plasmodium falciparum (p. fal.)* malaria parasites. The deep imputation model was able to impute this sparse and noisy matrix, while exploiting correlations between the strains, labs and associated measurements. The model also produced accurate error bars for both imputed results and virtual model predictions. A virtual library of approximately 1,000,000 structures was created through a combination of the two generative methods and the probabilistic scoring profile was defined to maximise activity in all assays, as well as increase solubility, while taking the confidence in each individual prediction into account.

Figure 2 shows the four compounds most likely to succeed, generated by the two approaches. Upon review, compound a) was considered to have a reactivity problem that may result in HF production and was discounted. In addition, compound c) was considered to be potentially unstable. The algorithms used could be improved to detect these high-level pitfalls and this reiterates the importance of the 'expert chemist verification' as depicted in Figure 1.

The compound most likely to succeed, the tert-butyl compound (b), had a predicted pIC_{50} value of 6.4 in the target assay. This was sent for experimental synthesis and, when tested, the experimental pIC_{50} was 6.2, well within the uncertainties in both the experimental and predicted values. This activity met the project criteria for activity, which was a $pIC_{50} > 6$.

Of the novel compounds proposed by four organisations in the project, this was the only experimentally verified active, illustrating that methods which cannot reliably consider uncertainty often struggle to filter the successful actives from noisy predictions. Furthermore, the project chemists considered that the tert-butoxy group on compound b) gives new directions in the SAR and it is unlikely that this compound would have been considered by a human. These are exactly

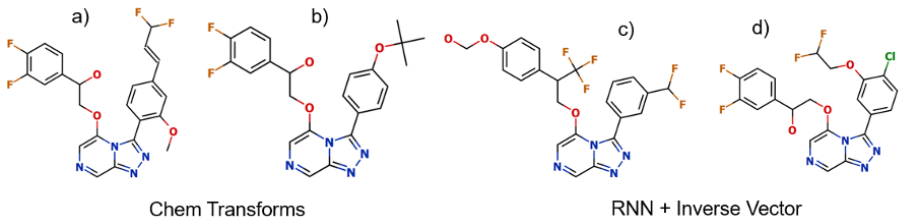

Figure 2: The most confident structures generated by the AI. Compounds a) and b) are from the bottom up Nova approach, compounds c) and d) were generated using the generative ML approach.

the kinds of benefits that an augmented tool will provide.

Additional insights were found in imputation of the original data in the context of the output 'confident identification of missed opportunities' in the process illustrated in Figure 1. Upon imputation, a compound with a single experimental "inactive" measurement (IC_{50} > 10) was proposed as being active in two other assays with a high degree of confidence (predicted *p. fal.* pEC_{50} of 7.2 and single shot inhibition of 96%). Upon digging deeper, the compound was found to be chiral (Figure 3). The compound was supposedly enantioenriched, but it was not known which enantiomer was more prevalent and the chirality had not been registered in the database. Upon experimentally resolving the enantiomers by chiral liquid chromatography, the enantiomer in Figure 3 was confirmed as inactive, while the other was confirmed as active. The AI was aware that there was a chance for activity and the active compound, which had been ruled out by a single datapoint, could have been a missed opportunity.

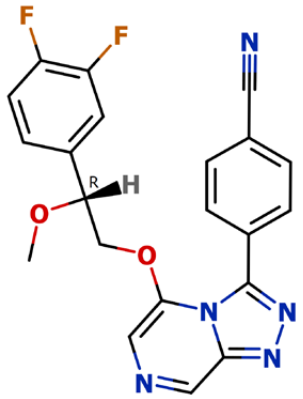

Figure 3: The chiral compound that was highlighted as a potential missed opportunity in the dataset.

Opportunities for Future Improvements

There are many further potential improvements than can be made to the process in Figure 1. We can combine additional data in the first instance to find correlations in public or private repositories. The descriptor generator used in the data preparation block could be upgraded to a machine learning tool such as a graph convolutional network. Further advances could be made to the optimisation steps, such as by employing reinforcement learning approaches. The RNN encoder could be replaced with a full generative adversarial network, which has been trained to produce drug-like and synthetically accessible compounds. The foundation developed here could easily grow into an even more powerful AI for drug discovery.

Conclusions

The combination of machine learning components can lead to an advanced system, which embodies a sophisticated AI tool. Using the Alchemite method, we are able to use all of the data available, even though it is sparse and noisy, and the resulting models output robust uncertainty estimates, which are essentially for later MPO and prioritisation of compounds. This concept was demonstrated through the efficient identification of a novel active antimalarial compound.

REFERENCES

1. Fenn, J.; Blosch, M. Understanding Gartner's Hype Cycles; 2018.
2. Davenport, T. H.; Glover, W. J. Artificial Intelligence and the Augmentation of Health Care Decision-Making. NEJM Catal. 2018, 4
3. Irwin, B. W. J.; Mahmoud, S.; Whitehead, T. M.; Conduit, G. J.; Segall, M. D. Imputation versus Prediction: Applications in Machine Learning for Drug Discovery. Futur. Drug Discov. 2020, 2 (2), FDD38.
4. Whitehead, T. M.; Irwin, B. W. J.; Hunt, P.; Segall, M. D.; Conduit, G. J. Imputation of Assay Bioactivity Data Using Deep Learning. J. Chem. Inf. Model. 2019, 59 (3), 1197–1204.
5. Irwin, B.W.J.; Levell, J.; Whitehead, T.; Segall, M.; Conduit, G. Practical Applications of Deep Learning to Impute Drug Discovery Data. J. Chem. Inf. Model. 2020.
6. Prykhodko, O.; Johansson, S. V.; Kotsias, P. C.; Arús-Pous, J.; Bjerrum, E. J.; Engkvist, O.; Chen, H. A de Novo Molecular Generation Method Using Latent Vector Based Generative Adversarial Network. J. Cheminform. 2019, 11 (1), 1–13.
7. Segall, M.; Champness, E.; Leeding, C.; Lilien, R.; Mettu, R.; Stevens, B. Applying Medicinal Chemistry Transformations and Multiparameter Optimization to Guide the Search for High-Quality Leads and Candidates. J. Chem. Inf. Model. 2011, 51 (11), 2967–2976.
8. D. Segall, M. Multi-Parameter Optimization: Identifying High Quality Compounds with a Balance of Properties. Curr. Pharm. Des. 2012, 18 (9), 1292–1310.
9. Segall, M. Advances in Multiparameter Optimization Methods for de Novo Drug Design. Expert Opin. Drug Discov. 2014, 9 (7), 803–817.
10. Segall, M. D.; Champness, E. J. The Challenges of Making Decisions Using Uncertain Data. J. Comput. Aided. Mol. Des. 2015, 29 (9), 809–816.
11. Swain, C.; Todd, M.; Kanza, S.; Frey, J. G. AI3SD, OSM & RSC-CICAG Predicting the Activity of Drug Candidates When There Is No Target Workshop Report; 2020.
12. Segall, M.; Mansley, T.; Hunt, P.; Champness, E. Capturing and Applying Knowledge to Guide Compound Optimisation. Drug Discov. Today 2019, 24 (5), 1074–1080.

Benedict Irwin

Benedict Irwin is a senior scientist at Optibrium Ltd based in the 'Silicon Fen' of Cambridge, UK. He has a background in theoretical physics and scientific computing and works on machine learning and mathematical methods for chemistry and structural biology problems. He has developed methods such as atom-wise free energy perturbation and `exact learning', and has recently worked on applying deep imputation to drug discovery problems.

Email: ben@optibrium.com

Alexander Wade

Alexander Wade is a PhD candidate in the Theory of Condensed Matter group at the University of Cambridge. Aiming to graduate in 2020, his work focuses on developing optimisation methods which can be applied to computational drug discovery using a range of methods including atomistic simulations and ligand-based machine learning models.

Email: adw62@cam.ac.uk

Matthew Segall

Matthew Segall is CEO of Optibrium and has an MSc in computation from the University of Oxford and a PhD in theoretical physics from the University of Cambridge. He has led teams developing predictive models and intuitive decision-support and visualisation tools for drug discovery. In 2009 he led a management buyout of the StarDrop business to found Optibrium, which develops software for small molecule design, optimisation, and data analysis.

Email: matt@optibrium.com