Published in final edited form as:

Nat Mater. 2019 May ; 18(5): 435–441. doi:10.1038/s41563-019-0338-z.

Exploiting Machine Learning for End-To-End Drug Discovery and Development

Sean Ekins¹, Ana C. Puhl¹, Kimberley M. Zorn¹, Thomas R. Lane¹, Daniel P. Russo^{1,2}, Jennifer J. Klein¹, Anthony J. Hickey^{3,4}, and Alex M. Clark⁵

¹Collaborations Pharmaceuticals, Inc. 840 Main Campus Drive, Lab 3510, Raleigh, NC 27606, USA.

²The Rutgers Center for Computational and Integrative Biology, Camden, NJ, 08102, USA.

³RTI International, Research Triangle Park, North Carolina, United States of America.

⁴UNC Catalyst for Rare Diseases, Eshelman School of Pharmacy, University of North Carolina at Chapel Hill, NC 27599-7536

⁵Molecular Materials Informatics, Inc., Montreal, Quebec, Canada.

Abstract

A variety of machine learning methods such as Naïve Bayesian, support vector machines and more recently deep neural networks are demonstrating their utility for drug discovery and development. These leverage the generally bigger data sets created from high throughput screening data and allow prediction of bioactivities for targets and molecular properties with increased levels of accuracy. We have only just begun to exploit the potential of these techniques but they may already be fundamentally changing the research process for identifying new molecules and/or repurposing old drugs. The integrated application of such machine learning models for end-to-end (E2E) application is broadly relevant and has considerable implications for developing future therapies and their targeting.

Learning from history

‘Those who do not remember the past are condemned to repeat it’ (Santayana). This observation applies as much to drug discovery as it does to other aspects of human endeavor¹. The history of drug discovery is a prelude to the emerging potential of computer-assisted data exploration. One constant in drug discovery is that every few years the estimated cost to develop drugs rises further. Less than 20 years ago, developing a drug took ~12 years, cost under a billion dollars, and the biggest challenges were failures due to efficacy or toxicity-induced attrition². *in vitro* pharmacological profiling implemented earlier in the drug discovery process helped to identify some predictable undesirable off-

*Correspondence to Sean Ekins sean@collaborationspharma.com.

Competing Financial Interests

SE is Founder and CEO, ACP, KMZ, TL and JJK are employees, DPR and AMC are consultants of Collaborations Pharmaceuticals, Inc. AMC is also the founder and owner of Molecular Materials Informatics, Inc. AJH has no conflicts of interest.

target activity profiles, which would hinder drug candidate development or even lead to market withdrawal if discovered after drug approval³. These technologies did not shorten the time or decrease the cost required to get a candidate drug to market, as now the costs of development are upwards of \$2.8 billion⁴. These methods have also not been able to predict clinical failures due to idiosyncratic toxicity and this may be due to the lack of *in vitro-in vivo* correlation⁵ while efficacy is as complex to predict⁶.

Today's pressures are likely different in each company, depending on the target market and available resources, although bottlenecks are common⁷. The biggest cost and most time-consuming component associated with drug development is conducting clinical trials. This is reflected in the high prices of these chronic treatments, which puts pressure on the United States healthcare and insurance systems as well as on the patients. Resolutions to these problems include efforts to speed up regulatory review and simplify clinical trials. Less than a decade ago one solution proposed was to increase the number and quality of innovative, cost-effective new medicines without incurring unsustainable research and development (R&D) costs⁸. The R&D process itself is recognized as far from the linear pathway commonly described (Fig 1). This is clarified in a Drug Discovery, Development, and Deployment Map (4DM)⁷. Another way to possibly improve productivity in this complex environment is to implement machine learning across all areas of drug discovery and development⁹ for which there is sufficient data to train models. Machine learning is a growing field of artificial intelligence that uses different statistical techniques to enable computers to learn from various data types without being explicitly programmed. This would be analogous to converting the 4DM from a 2D map into a functioning computational model that can be used to make predictions and requires radically re-thinking of the whole R&D process, learning from it and optimizing it (Fig 1). If models were available for all aspects of drug discovery and development, they could be used seamlessly to predict whether a compound was likely to be ultimately clinically viable (Fig 1). This process could be described as end-to-end (E2E). Potential limitations of a linear combination of models might appear as errors could accumulate depending on the accuracy of each model, which may then influence the overall utility and prediction. Other optimal combinations of models could also be developed as customized pipelines are developed depending on the disease, target and therapeutic type. These efforts build on the recent proposal that machine learning will impact the future of design, synthesis, characterization and application of molecules and materials¹⁰.

Many classification and clustering solutions in biology, medicine, precision phenotyping, and clinical diagnostic support systems have leveraged machine learning methods. A subset of these methods are "unsupervised learning" techniques that can be used to model and learn from multi-omics type data. Generically, this type of machine learning approach attempts to identify meaningful inferences from datasets that lack classification and categorical labels. For example, an approach called computational phenotyping has emerged to embrace the complexity inherent in disease mechanisms with machine learning to define accurate phenotypes and has been used to predict antibiotic resistance phenotypes in a variety of bacterial species¹¹. Rather than thinking of these different efforts in isolation we will need to integrate them in the complete discovery and development pathway.

Instead of proposing a single algorithm or approach as the optimal one, we subscribe to the concept that the limits of machine learning are likely to be exposed by experimental data inconsistency and dataset size, rather than the flaws of any individual modeling framework¹². The impact of this viewpoint is that machine learning can be applied for the same cost to identify treatments for a variety of diseases and may level the drug development field for smaller companies and researchers. Access to a collection of predictive models for many of the known diseases or related targets could help find additional compounds for testing that may have previously been overlooked¹³. Models for multiple targets already assist in the prediction of off-target effects¹⁴ as well as predicting the most potent compound-target interactions¹⁵.

The tipping point for machine learning

In order to build machine learning models high quality data are needed. The last ten years have seen a dramatic increase in the amount of public chemical and biological data in PubChem¹⁶, ChEMBL¹⁷, and other databases that include screening data. We now have millions of molecules and bioactivities for different disease targets as well as for absorption, distribution, metabolism, excretion and toxicology (ADME/Tox) properties. These data are an extremely valuable resource for drug discovery machine learning applications. We would suggest for many of the targets and diseases we now have plentiful data as indicated by the wide use of ChEMBL and the accuracy of the models generated¹⁸. A considerable limitation of many databases is the data are not 'model-ready' or machine-readable¹⁹ which is needed to successfully use any machine learning methodology. When you can download data in electronic formats, expert curation is always required to ensure compatibility of data and domain expertise is important to build and use machine learning models. ChEMBL does a much better job of curation of data but it still takes some effort to prepare for building a machine learning model.

Machine learning models such as support vector machines²⁰, k-Nearest Neighbors²¹, Naïve Bayes²², Random Forest²³ and many other methods²⁴, have long been utilized for drug discovery. However, recent interest in deep learning or deep neural networks (DNNs) for drug discovery has catalyzed interest in machine learning in this field more broadly. DNNs have been used in pattern recognition and machine learning²⁵, sparking their use in pharmacology and drug discovery²⁶ and becoming a source for numerous recent reviews¹². DNNs have been used in various pharmaceutical applications from docking to virtual screening and beyond (Table 1), but the rise in prominence is linked to increased computational power and the availability of larger datasets. While DNNs are inspired by biological neural networks and consist of layers of interconnected neurons; much of the interest in them is centered around the flexibility of their architecture²⁷ which allows the generation of models for single task or multitask machine learning²⁸ as well as predicting drug-target interactions²⁹. The use of DNNs is still in its relative infancy and has limited applications for cheminformatics as compared with other methods³⁰. While DNN algorithms are increasingly available⁹, they are not 'plug and play' and their use takes significant time to optimize. Also, the selection of which machine learning algorithm to use with each dataset is not readily predictable and there is really no agreement as to which algorithm is the best for cheminformatics versus other uses. One group has suggested using several

“benchmark” datasets for comparing the predictive ability of different molecular machine learning algorithms³¹, while others have performed comparisons using target related datasets from ChEMBL¹⁸. The assessments for DNNs could also be applied to essentially any public drug discovery relevant small molecule bioactivity dataset³², but to date this algorithm has rarely been used for prospective prediction and in some respects this is a limitation of many of the machine learning drug discovery studies published. DNNs have also been used to create novel features / descriptors from their molecular structure as an alternative to traditional molecular descriptors³³. 2D structures of molecules have seen use as an input to predict toxicity using the Tox21 benchmark set³⁴, which is also part of a platform called MoleculeNet³⁵. Generative DNNs have also been described for the generation virtual libraries of molecules and these enable *de novo* drug design with optimized properties³⁶. However, it should be pointed out that such proof of concept studies have not synthesized molecules to validate the predictions and this needs to happen to provide evidence of their value. The closest example to this ideal scenario has purchased close analogs to the molecules generated with generative DNNs and tested them against different kinases at 10 μ M, identifying several hits³⁷.

While there are several machine learning frameworks and tools available today aimed at using small molecules and related data, they are therefore not at the point where they are universally accessible to all scientists. The requirement for expert users in many ways has been the Achilles’ heel of cheminformatics, whereas computational tools for bioinformatics have found broader use due to their accessibility. Therefore, we need to rethink how to generally make the machine learning models for pharma more usable and user-friendly to increase the number of potential users and applications possible. We are at a clear tipping point for machine learning and deep learning in particular, but it has taken decades to reach this point and yet full integration of these models is still likely to be a work in progress.

Machine learning models in action

Machine learning methods in the pharmaceutical industry are most commonly used for virtual screening of compounds, reducing the need to generate more high-throughput screening data by cherry-picking compounds and performing low to medium-throughput screening³⁸. The same machine learning algorithms have been used widely in both pharmaceutical and toxicological research³⁰ (Table 1). Statistical machine learning methods have also been used to interrogate, model, and learn from complex multi-omics data to help to address uncertainties about the connections between different types of data³⁹. For example, machine learning methods have been applied to electronic health records to accurately predict multiple medical events from different centers without site-specific data harmonization, with recent data suggesting that deep learning was comparable to regularized logistic regression in this case⁴⁰.

Several of our own recent cheminformatics prospective testing efforts have identified compounds active *in vitro* and *in vivo* against Chagas disease¹³ and the Ebola virus⁴¹ using Bayesian algorithms. This Bayesian approach has also been widely applied to ADME properties by predicting aqueous solubility, mouse liver microsomal stability⁴², Caco-2 cell permeability⁴³, cytotoxicity⁴⁴ and interactions with transporters⁴⁵. We have also used many

different machine learning algorithms and descriptors in parallel to identify the optimum combination⁴⁶ and address complex problems facing the pharmaceutical industry related to the challenges of improving solubility or metabolic stability⁴⁷ while retaining bioactivity. These challenges still persist partly because the datasets may not cover sufficient chemical space, and the test molecules could be outside the applicability of the training set of the models. Optimization and understanding the application of machine learning models is generally not trivial. Instead the field has tended to emphasize discovering the ‘perfect individual model’ and using various forms of cross validation to evaluate them.

We and others have recently performed several analyses using diverse drug discovery datasets and metrics to compare different machine learning methods using one type of frequently used molecular descriptor, namely FCFP6 fingerprints⁴⁸. After 4-fold cross validation and ranked normalized scores of metrics, DNNs ranked higher than all the other machine learning methods across all datasets⁴⁸. Other researchers have also compared several machine learning approaches with different datasets from ChEMBL using random split and temporal cross validation to show the superiority of DNNs⁴⁹, or 5-fold cross validation and leave out 40% as a validation set⁵⁰. A nested cluster-cross validation strategy has also been used to show that DNNs outperform these other machine learning methods¹⁸. We followed these studies by assessing different machine learning methods and molecular descriptors with 18,886 compounds screened against *Mycobacterium tuberculosis*⁵¹. This comparison demonstrated that DNNs and support vector machines appear to be superior methods regardless of the descriptor type for training and 5-fold cross validation. Conversely, external testing of DNN models with a large test set did not perform as well as other machine learning methods. More recently we have evaluated these same machine learning algorithms and descriptors for multiple estrogen receptor datasets⁴⁶. For predicting compounds within the training set, DNNs had higher accuracy than other methods in 5-fold cross validation. For external test set predictions DNN and most classic machine learning models perform similarly regardless of dataset or molecular descriptors⁴⁶. The fact that DNN does not always outperform other methods for external testing and in several cases is not the best, is important to consider due to the computational cost of DNN. This therefore deserves more exhaustive assessment to determine which algorithm to use with each dataset. Our own efforts continue to reflect this pattern, namely that while DNN excels at cross validation assessments it is generally no better than other machine learning methods for external testing.

Models for all diseases

The majority of global pharmaceutical companies are focused on the major diseases (e.g. cancer, cardiovascular, pain, diabetes, arthritis) that conform to a robust business model, and most research scientists are similarly engaged in these endeavors. However, other diseases that in aggregate involve large patient populations and represent major unmet medical needs are gaining attention. There are neglected and tropical diseases (e.g. malaria, tuberculosis and others) and rare diseases (defined as affecting less than 200,000 people in the United States) in which interest has increased as the FDA priority voucher⁵² has provided an incentive for companies to develop new treatments. While we have focused our efforts on neglected and tropical disease machine learning models for tuberculosis and malaria, these

diseases are in the enviable position of having very large datasets (>300,000 compounds) from high-throughput screening which can be utilized for machine learning⁵³. While NIH funding has gone into the high throughput screening⁵⁴, until recently there have not been comparable efforts on data mining or machine learning with this data¹⁶. Recently, we have stressed the need to scale and “industrialize” rare disease drug discovery⁵⁵ and move towards higher throughput and collaborative approaches. We have also proposed that machine learning could be used to find treatments for rare diseases using an iterative approach⁵⁶. This methodology would involve first linking the targets for rare diseases, building models for targets related to these diseases, and then use machine learning to identify additional molecules for future testing and validation (Fig 2). Currently available chemical and biological data relevant for rare disease drug discovery is available but diffuse, existing in an array of public or private databases. Several recent efforts have focused on developing pipelines using natural language processing and human curation to mine promising targets for drug development for rare diseases⁵⁷. These examined diseases with late onset, but clearly there is also an urgent need to address rare diseases with an early onset. Others have initiated different approaches to combat rare diseases through the development of a comprehensive global genotype-phenotype database⁵⁸, sharing genomics data⁵⁹ or other aspects of rare diseases⁶⁰, as well as assist patients and caregivers⁶¹. To date there are no efforts specifically using machine learning to identify drugs for rare diseases that leverage the relevant datasets for targets that are in the public domain. These approaches could learn from the work performed on neglected and tropical diseases which has used public datasets to identify new compounds⁵⁴.

Making models more accessible and interpretable

If we can combine high-quality curated screening data with cutting edge machine learning algorithms and molecular descriptors, there is the opportunity to build models that can be used to reliably predict new molecules for most areas of drug discovery. An iterative loop can be created where a group leverages its expertise with data and models to propose molecules, the experimentalists procure them and measure bioactivities, and the results are returned in a form that can be inserted directly into the model building process (Fig 2). This type of simplistic approach while limited to drug discovery is amenable to both large and small company efforts and can be used across many projects simultaneously to create a pipeline of internal and external projects. We have taken this strategy with our own machine learning software called Assay Central⁴⁶ (Fig 2) representing an accessible approach to scale drug discovery⁸. There is now increasing focus on machine learning in drug discovery which suggests the utility of such approaches is increasing. This is exemplified by the number of deals between start-up machine learning-based drug discovery companies and big pharma, biotechs, or VC investors⁶². While making machine learning models and predictions more accessible is important to demonstrate impact, efforts to increase the interpretability of these models beyond the “black box” are critical⁶³. We and others have taken different routes to improve this aspect including tools to highlight contributions of models to test molecules⁶⁴, identifying training compounds in the same neighborhood as test molecules and scores of model applicability or overlap^{46, 63}.

Nanoparticles and nanomedicines

Machine learning may be applicable to developing nanomedicines by exploiting large datasets, in an analogous manner to other areas of drug discovery³⁴. These efforts can enable the quantitative prediction of desirable molecules before synthesis and focus research on experiments with the most promising candidates. The field of nanomedicine has led to the development of nanoinformatics and the use of data mining and machine learning to develop nano-QSARs to predict functional and structural properties of nanoparticles. A relatively wide array of machine learning approaches³⁴ have been applied to prediction of different biomedical properties of nanoparticles such as predicting cellular uptake, cytotoxicity, molecular loading, molecular release, nanoparticle adherence, nanoparticle size and polydispersity⁶⁵. Computational methods can be used to predict the particle self-assembly process for targeted drug carrier nanoparticles. Quantitative structure-nanoparticle assembly prediction models have been used to generate predictions of nano-assembly which were found to encapsulate drugs with high loadings and have then also been validated in cancer models⁶⁶. Interestingly, a machine learning method has also been described to identify clinical trials involving nanodrugs and nanodevices from [ClinicalTrials.gov](https://clinicaltrials.gov)⁶⁷. While drug discovery has seen decades of applications of machine learning, for nanoparticle research there are far fewer examples and data available for model building are limited to select nanomaterial databases⁶⁸. As most of the published examples use small datasets (10–100s of molecules), deep learning has limited value and has rarely been applied. Commercially available tools that could enable scientists to develop nanomedicines using these models have yet to be developed to date. The relevance of nanomaterial-related approaches is in drug formulation or delivery and should be considered an integral part of E2E.

Machine Learning repurposing

Our thinking need not be limited to discovering just new molecular entities as machine learning can help explore the patterns of known drugs and their interactions with drug targets and potentially repurpose already approved molecules. Much of this information on potential repurposing opportunities can already be gleaned from public sources⁶⁹. This suggests that we may not even need to screen large numbers of compounds in future. We now have an abundance of data, powerful and plentiful computing, public and private efforts to develop databases and models and accessible ways to test compounds on a fee-for-service basis. Multiscale models defining networks for a given disease can also be used to construct gene expression assays for high-throughput screening. While these are relatively nascent as we learn more about biology their impact will also expand. For example, in a classic paradigm inflammatory bowel disease (IBD) signatures were derived from surgical specimens and intersected with Connectivity Map⁷⁰ data representing transcriptional readouts across a number of cell lines in response to treatment with many hundreds of drugs using a novel pattern-matching algorithm⁷¹. From this research the anticonvulsant drug topiramate was identified and experimentally validated as a novel treatment for IBD. The same approach has been applied to transcriptional profiles of non-small cell lung cancer (NSCLC) which identified imipramine, bepridil, and promethazine and cimetidine as NSCLC inhibitors⁷². The increasing number of drug-repositioning investigations, suggests that the reuse of medications for common, rare or orphan diseases is a viable approach²⁶.

Machine learning can help to assess whether a drug can be repositioned for a novel indication⁷³ There are likely many examples of approved drugs finding additional uses for various diseases⁵⁴ and obviously the key is to ensure that these reach the patient in a timely manner.

Repurposing efforts could be greatly assisted by obtaining data from inside companies and academic research institutes. Often a major concern about making predictions with computational models is the vastness of chemical space and the potential for selecting compounds with unfavorable properties. These limitations would be mitigated to some extent if previously approved drugs that have extensive ADMET data (or better still have reached phase I trials) could be repurposed.

The complete E2E model

In summary, while much of the focus of this review has been on cheminformatics, we propose that many areas across the pharmaceutical R&D spectrum and outside of it are ripe for machine learning (Table 1). Machine learning can learn from almost any data type, such as that from research papers, patient records, images, genes, symptoms, diseases, proteins, tissues, species and drug candidates or compounds that have been shown to affect any of the preceding⁷⁴. We could also imagine a complex interaction network between proteins upstream and downstream in the pathway that might dictate if the drug/s will work. Proteins have isoforms and redundancy, thus, inhibition of one might not be enough to illicit the desired response. In the same way, inhibition of one pathway might not be enough to achieve the response, since the cell has other pathway mechanisms that would be activated to circumvent the one that has been inhibited. In this context, we can apply machine learning to the whole pathway to evaluate how a network of protein interactions will react to a perturbation in the system such as the drug that is acting on a particular target and this in turn could lead to more personalized medicine⁷⁵. Rather than repeating the mistakes of the past, it is necessary to understand the biological context that gives rise to the disease and which gene network and proteins are operating before beginning drug discovery screens.

Using machine learning to integrate diverse, large-scale data can provide a path to predict which drug effects might best counteract the molecular networks underlying disease or result in less toxicity. This leads us to selecting the best targets and may ultimately help us to predict efficacy. Some approaches using machine learning methods have been developed to detect drug-target interactions⁷⁶, which is fundamental to both new drug discovery and drug repositioning. CRISPR-Cas9, has the potential to edit and renovate the harmful genes for personalized therapy and machine learning methods have been applied to predict the off-targets of CRISPR-Cas9 gene editing⁷⁷. Recently, a Cancer Drug Response profile scan (CDRscan) was developed, that predicts somatic mutation profile-based drug responsiveness by linking the tumor genomic fingerprint and its sensitivity to drugs and identified 14 oncology and 23 non-oncology drugs as having new potential cancer indications which may result in treatments tailored for each individual patient⁷⁸. In the area of drug safety, a random forest classifier was used to predict the effect of drugs on the fetus. The models successfully identified category C drugs that are likely to be harmful and those likely to be safe for fetal loss or congenital anomalies⁷⁹.

The future drug discovery and development process will use machine learning E2E which will impact training of the workforce. There will be a much heavier computational emphasis as they manage a dashboard of projects, molecules, and targets across all the aspects of the process and outcomes are predicted in parallel (Fig 1). Small pharmaceutical companies may then be able to address tens to hundreds of diseases computationally before narrowing to the most promising projects based on a wide variety of these computational models. Using machine learning more broadly across the industry could allow us to move beyond the limitations defined by researcher specialty and data silos but it will be important to perform prospective validation of the models to demonstrate progress. These efforts will increasingly demonstrate that machine learning algorithms can help us to discover the next generation of drugs.

Acknowledgements

In memory of Rebecca J. Williams.

Dr. Joel Freundlich, Dr. Renée J.G. Arnold, Dr. Peter Madrid, Dr. Jair Lage de Siqueira-Neto, Dr. Antony Williams, Dr. Alex Tropsha, Dr. Aaron Gerlach, Mr. Jacob Gerlach, Ms. Debra Chipman, Ms. Audrey Davidow and Dr. Maggie Hupcey are kindly acknowledged for discussions and some of the collaborations described herein.

SE acknowledges funding to Collaborations Pharmaceuticals Inc. from NIGMS R44 GM122196-02A1, NINDS 1R43NS107079-01, NINDS 3R43NS107079-01S1, NCATS 1UH2TR002084-01 and FY2018 UNC Research Opportunities Initiative (ROI) Award.

Research reported in this publication was supported by the National Institute Of Neurological Disorders And Stroke of the National Institutes of Health under Award Number R43NS107079. The content is solely the responsibility of the authors and does not necessarily represent the official views of the National Institutes of Health."

References

1. Butler LD, et al. Current nonclinical testing paradigms in support of safe clinical trials: An IQ Consortium DruSafe perspective. *Regul Toxicol Pharmacol*, 87 Suppl 3, S1–S15 (2017).
2. Kola I& Landis J Can the Pharmaceutical industry reduce attrition rates. *Nat Rev Drug Discov*, 3, 711–715 (2004). [PubMed: 15286737]
3. Bowes J, et al. Reducing safety-related drug attrition: the use of in vitro pharmacological profiling. *Nat Rev Drug Discov*, 11, 909–922 (2012). [PubMed: 23197038]
4. DiMasi JA, Grabowski HG& Hansen RW Innovation in the pharmaceutical industry: New estimates of R&D costs. *J Health Econ*, 47, 20–33 (2016). [PubMed: 26928437]
5. Kenna JG Human biology-based drug safety evaluation: scientific rationale, current status and future challenges. *Expert Opin Drug Metab Toxicol*, 13, 567–574 (2017). [PubMed: 28150517]
6. Gayvert KM, Madhukar NS& Elemento OA Data-Driven Approach to Predicting Successes and Failures of Clinical Trials. *Cell Chem Biol*, 23, 1294–1301 (2016). [PubMed: 27642066]
7. Wagner JA, et al. Application of a Dynamic Map for Learning, Communicating, Navigating, and Improving Therapeutic Development. *Clin Transl Sci*, 11, 166–174 (2018). [PubMed: 29271559]
8. Paul SM, et al. How to improve R&D productivity: the pharmaceutical industry's grand challenge. *Nat Rev Drug Discov*, 9, 203–214 (2010). [PubMed: 20168317]
9. Zhavoronkov A Artificial Intelligence for Drug Discovery, Biomarker Development, and Generation of Novel Chemistry. *Mol Pharm*, 15, 4311–4313 (2018). [PubMed: 30269508]
10. Davies DW, Butler KT, Isayev O& Walsh A Materials discovery by chemical analogy: role of oxidation states in structure prediction. *Faraday Discuss*, 211, 553–568 (2018). [PubMed: 30027179]
11. Drouin A, et al. Predictive computational phenotyping and biomarker discovery using reference-free genome comparisons. *BMC Genomics*, 17, 754 (2016). [PubMed: 27671088]

12. Chen H, et al. The rise of deep learning in drug discovery. *Drug Discov Today*, 23, 1241–1250 (2018). [PubMed: 29366762]
13. Ekins S, et al. Machine Learning Models and Pathway Genome Data Base for *Trypanosoma cruzi* Drug Discovery *PLoS Negl Trop Dis*, 9, e0003878 (2015). [PubMed: 26114876]
14. Lampa S, et al. Predicting Off-Target Binding Profiles With Confidence Using Conformal Prediction. *Front Pharmacol*, 9, 1256 (2018). [PubMed: 30459617]
15. Reker D, Rodrigues T, Schneider P & Schneider G Identifying the macromolecular targets of de novo-designed chemical entities through self-organizing map consensus. *Proc Natl Acad Sci U S A*, 111, 4067–4072 (2014). [PubMed: 24591595]
16. Kim S, et al. PubChem Substance and Compound databases. *Nucleic Acids Res*, 44, D1202–1213 (2016). [PubMed: 26400175]
17. Gaulton A, et al. ChEMBL: a large-scale bioactivity database for drug discovery. *Nucleic Acids Res*, 40, D1100–1107 (2012). [PubMed: 21948594]
18. Mayr A, et al. Large-scale comparison of machine learning methods for drug target prediction on ChEMBL. *Chem Sci*, 9, 5441–5451 (2018). [PubMed: 30155234]
19. Clark AM, Williams AJ & Ekins S Machines first, humans second: on the importance of algorithmic interpretation of open chemistry data. *J Cheminform* 7, 9 (2015). [PubMed: 25798198]
20. Christianini N & Shawe-Taylor J Support vector machines and other kernel-based learning methods. (Cambridge University Press: Cambridge, MA, 2000).
21. Shen M, et al. Development and validation of k-nearest neighbour QSPR models of metabolic stability of drug candidates. *J Med Chem*, 46, 3013–3020 (2003). [PubMed: 12825940]
22. Bender A, et al. Analysis of Pharmacology Data and the Prediction of Adverse Drug Reactions and Off-Target Effects from Chemical Structure. *ChemMedChem*, 2, 861–873 (2007). [PubMed: 17477341]
23. Susnow RG & Dixon SL Use of robust classification techniques for the prediction of human cytochrome P450 2D6 inhibition. *J Chem Inf Comput Sci*, 43, 1308–1315 (2003). [PubMed: 12870924]
24. Mitchell JB Machine learning methods in chemoinformatics. *Wiley Interdiscip Rev Comput Mol Sci*, 4, 468–481 (2014). [PubMed: 25285160]
25. Schmidhuber J Deep learning in neural networks: an overview. *Neural Netw*, 61, 85–117 (2015). [PubMed: 25462637]
26. Aliper A, et al. Deep Learning Applications for Predicting Pharmacological Properties of Drugs and Drug Repurposing Using Transcriptomic Data. *Mol Pharm*, 13, 2524–2530 (2016). [PubMed: 27200455]
27. Ma J, et al. Deep neural nets as a method for quantitative structure-activity relationships. *J Chem Inf Model*, 55, 263–274 (2015). [PubMed: 25635324]
28. Wu K, Zhao Z, Wang R & Wei GW TopP-S: Persistent homology-based multi-task deep neural networks for simultaneous predictions of partition coefficient and aqueous solubility. *J Comput Chem* (2018).
29. Wen M, et al. Deep-Learning-Based Drug-Target Interaction Prediction. *J Proteome Res*, 16, 1401–1409 (2017). [PubMed: 28264154]
30. Ekins S The next era: Deep learning in pharmaceutical research. *Pharm Res*, 33, 2594–2603 (2016). [PubMed: 27599991]
31. Wu Z, et al. MoleculeNet: A Benchmark for Molecular Machine Learning. 2016 [cited] Available from: <https://arxiv.org/ftp/arxiv/papers/1703/1703.00564.pdf>
32. Altae-Tran H, Ramsundar B, Pappu AS & Pande V Low Data Drug Discovery with One-Shot Learning. *ACS Cent Sci*, 3, 283–293 (2017). [PubMed: 28470045]
33. Kadurin A, et al. druGAN: An Advanced Generative Adversarial Autoencoder Model for de Novo Generation of New Molecules with Desired Molecular Properties in Silico. *Mol Pharm*, 14, 3098–3104 (2017). [PubMed: 28703000]
34. Butler KT, et al. Machine learning for molecular and materials science. *Nature*, 559, 547–555 (2018). [PubMed: 30046072]

35. Rifaioğlu AS, et al. Recent applications of deep learning and machine intelligence on in silico drug discovery: methods, tools and databases. *Brief Bioinform* (2018).
36. Popova M, Isayev O& Tropsha A Deep reinforcement learning for de novo drug design. *Sci Adv*, 4, eaap7885 (2018). [PubMed: 30050984]
37. Putin E, et al. Adversarial Threshold Neural Computer for Molecular de Novo Design. *Mol Pharm*, 15, 4386–4397 (2018). [PubMed: 29569445]
38. McGaughey GB, et al. Comparison of topological, shape, and docking methods in virtual screening. *J Chem Inf Model*, 47, 1504–1519 (2007). [PubMed: 17591764]
39. Johnson KW, et al. Enabling Precision Cardiology Through Multiscale Biology and Systems Medicine. *JACC Basic Transl Sci*, 2, 311–327 (2017). [PubMed: 30062151]
40. Rajkomar A, et al. Scalable and accurate deep learning with electronic health records. *Digital Medicine*, 1, 18 (2018).
41. Ekins S, et al. Machine learning models identify molecules active against Ebola virus in vitro. *F1000Res*, 4, 1091 (2015). [PubMed: 26834994]
42. Perryman AL, Stratton TP, Ekins S& Freundlich JS Predicting mouse liver microsomal stability with “pruned” machine learning models and public data. *Pharm Res*, 33, 433–449 (2015). [PubMed: 26415647]
43. Clark AM, et al. Open source bayesian models: 1. Application to ADME/Tox and drug discovery datasets. *J Chem Inf Model*, 55, 1231–1245 (2015). [PubMed: 25994950]
44. Perryman AL, et al. Naive Bayesian Models for Vero Cell Cytotoxicity. *Pharm Res*, 35, 170 (2018). [PubMed: 29959603]
45. Sandoval PJ, et al. Assessment of Substrate Dependent Ligand Interactions at the Organic Cation Transporter OCT2 Using Six Model Substrates. *Mol Pharmacol*, 94, 1057–1068 (2018). [PubMed: 29884691]
46. Russo DP, et al. Comparing Multiple Machine Learning Algorithms and Metrics for Estrogen Receptor Binding Prediction. *Mol Pharm*, 15, 4361–4370 (2018). [PubMed: 30114914]
47. Stratton TP, et al. Addressing the Metabolic Stability of Antituberculars through Machine Learning. *ACS Med Chem Lett*, 8, 1099–1104 (2017). [PubMed: 29057058]
48. Korotcov A, Tkachenko V, Russo DP& Ekins S Comparison of Deep Learning With Multiple Machine Learning Methods and Metrics Using Diverse Drug Discovery Datasets. *Mol Pharm*, 14, 4462–4475 (2018).
49. Lenselink EB, et al. Beyond the hype: deep neural networks outperform established methods using a ChEMBL bioactivity benchmark set. *J Cheminform*, 9, 45 (2017). [PubMed: 29086168]
50. Koutsoukas A, Monaghan KJ, Li X& Huan J Deep-learning: investigating deep neural networks hyper-parameters and comparison of performance to shallow methods for modeling bioactivity data. *J Cheminform*, 9, 42 (2017). [PubMed: 29086090]
51. Lane T, et al. Comparing and Validating Machine Learning Models for Mycobacterium tuberculosis Drug Discovery. *Mol Pharm* (2018).
52. Ridley DB Priorities for the Priority Review Voucher. *Am J Trop Med Hyg*, 96, 14–15 (2017). [PubMed: 27573624]
53. Ekins S, et al. Bayesian Models Leveraging Bioactivity and Cytotoxicity Information for Drug Discovery. *Chem Biol*, 20, 370–378 (2013). [PubMed: 23521795]
54. Hernandez HW, et al. High Throughput and Computational Repurposing for Neglected Diseases *Pharm Res*, 36, (2018).
55. Ekins S Industrializing rare disease therapy discovery and development. *Nat Biotechnol*, 35, 117–118 (2017). [PubMed: 28178258]
56. Ekins S& Perlstein EO Doing it All - How Families are Reshaping Rare Disease Research. *Pharm Res*, 35, 1–4 (2018).
57. Chen B& Altman RB Opportunities for developing therapies for rare genetic diseases: focus on gain-of-function and allostery. *Orphanet J Rare Dis*, 12, 61 (2017). [PubMed: 28412959]
58. Trujillano D, et al. A comprehensive global genotype-phenotype database for rare diseases. *Mol Genet Genomic Med*, 5, 66–75 (2017). [PubMed: 28116331]

59. Thompson R, et al. RD-Connect: an integrated platform connecting databases, registries, biobanks and clinical bioinformatics for rare disease research. *J Gen Intern Med*, 29 Suppl 3, S780–787 (2014). [PubMed: 25029978]
60. Rath A, et al. Representation of rare diseases in health information systems: the Orphanet approach to serve a wide range of end users. *Hum Mutat*, 33, 803–808 (2012). [PubMed: 22422702]
61. Anon. Rare Disease InfoHub. 2018 [cited]Available from: <https://rarediseases.oscar.ncsu.edu/>
62. Fleming N How artificial intelligence is changing drug discovery. *Nature*, 557, S55–S57 (2018). [PubMed: 29849160]
63. Chuang KV& Keiser MJ Adversarial Controls for Scientific Machine Learning. *ACS Chem Biol*, 13, 2819–2821 (2018). [PubMed: 30336670]
64. Marchese Robinson RL, Palczewska A, Palczewski J& Kidley N Comparison of the Predictive Performance and Interpretability of Random Forest and Linear Models on Benchmark Data Sets. *J Chem Inf Model*, 57, 1773–1792 (2017). [PubMed: 28715209]
65. Jones DE, Ghandehari H& Facelli JC A review of the applications of data mining and machine learning for the prediction of biomedical properties of nanoparticles. *Comput Methods Programs Biomed*, 132, 93–103 (2016). [PubMed: 27282231]
66. Shamay Y, et al. Quantitative self-assembly prediction yields targeted nanomedicines. *Nat Mater*, 17, 361–368 (2018). [PubMed: 29403054]
67. de la Iglesia D, et al. A machine learning approach to identify clinical trials involving nanodrugs and nanodevices from [ClinicalTrials.gov](https://clinicaltrials.gov) PLoS One, 9, e110331 (2014). [PubMed: 25347075]
68. Tropsha A, Mills KC& Hickey AJ Reproducibility, sharing and progress in nanomaterial databases. *Nat Nanotechnol*, 12, 1111–1114 (2017). [PubMed: 29209007]
69. Baker NC, Ekins S, Williams AJ& Tropsha A A bibliometric review of drug repurposing. *Drug Discov Today*, 23, 661–672 (2018). [PubMed: 29330123]
70. Lamb J, et al. The Connectivity Map: using gene-expression signatures to connect small molecules, genes, and disease. *Science*, 313, 1929–1935 (2006). [PubMed: 17008526]
71. Dudley JT, et al. Computational repositioning of the anticonvulsant topiramate for inflammatory bowel disease. *Sci Transl Med*, 3, 96ra76 (2011).
72. Schadt EE, Buchanan S, Brennand KJ& Merchant KM Evolving toward a human-cell based and multiscale approach to drug discovery for CNS disorders. *Front Pharmacol*, 5, 252 (2014). [PubMed: 25520658]
73. Napolitano F, et al. Drug repositioning: a machine-learning approach through data integration. *J Cheminform*, 5, 30 (2013). [PubMed: 23800010]
74. Cruz S, et al. In Silico HCT116 Human Colon Cancer Cell-Based Models En Route to the Discovery of Lead-Like Anticancer Drugs. *Biomolecules*, 8, (2018).
75. Frohlich H, et al. From hype to reality: data science enabling personalized medicine. *BMC Med*, 16, 150 (2018). [PubMed: 30145981]
76. Chen R, et al. Machine Learning for Drug-Target Interaction Prediction. *Molecules*, 23, (2018).
77. Lin J& Wong KC Off-target predictions in CRISPR-Cas9 gene editing using deep learning. *Bioinformatics*, 34, i656–i663 (2018). [PubMed: 30423072]
78. Chang Y, et al. Cancer Drug Response Profile scan (CDRscan): A Deep Learning Model That Predicts Drug Effectiveness from Cancer Genomic Signature. *Sci Rep*, 8, 8857 (2018). [PubMed: 29891981]
79. Boland MR, Polubriaginof F& Tatonetti NP Development of A Machine Learning Algorithm to Classify Drugs Of Unknown Fetal Effect. *Sci Rep*, 7, 12839 (2017). [PubMed: 28993650]
80. Zang Q, et al. In Silico Prediction of Physicochemical Properties of Environmental Chemicals Using Molecular Fingerprints and Machine Learning. *J Chem Inf Model*, 57, 36–49 (2017). [PubMed: 28006899]
81. Lusci A, Pollastri G& Baldi P Deep architectures and deep learning in chemoinformatics: the prediction of aqueous solubility for drug-like molecules. *J Chem Inf Model*, 53, 1563–1575 (2013). [PubMed: 23795551]

82. Hong H, Thakkar S, Chen M& Tong W Development of Decision Forest Models for Prediction of Drug-Induced Liver Injury in Humans Using A Large Set of FDA-approved Drugs. *Sci Rep*, 7, 17311 (2017). [PubMed: 29229971]
83. Korotcov A, Tkachenko V, Russo DP& Ekins S Comparison of Deep Learning With Multiple Machine Learning Methods and Metrics Using Diverse Drug Discovery Data Sets. *Mol Pharm*, 14, 4462–4475 (2017). [PubMed: 29096442]
84. Wang W, Kim MT, Sedykh A& Zhu H Developing Enhanced Blood-Brain Barrier Permeability Models: Integrating External Bio-Assay Data in QSAR Modeling. *Pharm Res*, 32, 3055–3065 (2015). [PubMed: 25862462]
85. Baba H, Takahara J, Yamashita F& Hashida M Modeling and Prediction of Solvent Effect on Human Skin Permeability using Support Vector Regression and Random Forest. *Pharm Res*, 32, 3604–3617 (2015). [PubMed: 26033768]
86. Xu C, et al. In silico prediction of chemical Ames mutagenicity. *J Chem Inf Model*, 52, 2840–2847 (2012). [PubMed: 23030379]
87. Huang W, et al. Prediction of human clearance based on animal data and molecular properties. *Chem Biol Drug Des*, 86, 990–997 (2015). [PubMed: 25845625]
88. Basant N, Gupta S& Singh KP QSAR modeling for predicting reproductive toxicity of chemicals in rats for regulatory purposes. *Toxicol Res (Camb)*, 5, 1029–1038 (2016). [PubMed: 30090410]
89. Alhalaweh A, et al. Computational predictions of glass-forming ability and crystallization tendency of drug molecules. *Mol Pharm*, 11, 3123–3132 (2014). [PubMed: 25014125]
90. Miller TH, et al. Prediction of bioconcentration factors in fish and invertebrates using machine learning. *Sci Total Environ*, 648, 80–89 (2019). [PubMed: 30114591]
91. Rose S, Bergquist SL& Layton TJ Computational health economics for identification of unprofitable health care enrollees. *Biostatistics*, 18, 682–694 (2017). [PubMed: 28369273]
92. Calderon CP, Daniels AL& Randolph TW Deep Convolutional Neural Network Analysis of Flow Imaging Microscopy Data to Classify Subvisible Particles in Protein Formulations. *J Pharm Sci*, 107, 999–1008 (2018). [PubMed: 29269269]
93. Degardin K, Guillemain A, Guerreiro NV& Roggo Y Near infrared spectroscopy for counterfeit detection using a large database of pharmaceutical tablets. *J Pharm Biomed Anal*, 128, 89–97 (2016). [PubMed: 27236101]
94. Page D, et al. Identifying Adverse Drug Events by Relational Learning. *Proc Conf AAAI Artif Intell*, 2012, 790–793 (2012). [PubMed: 24955289]
95. Rannals MD, et al. Psychiatric Risk Gene Transcription Factor 4 Regulates Intrinsic Excitability of Prefrontal Neurons via Repression of SCN10a and KCNQ1. *Neuron*, 90, 43–55 (2016). [PubMed: 26971948]

Learning and Process Optimization

Small Molecule Drug Discovery and Development Timeline

Figure 1. Implementing end-to-end (E2E) machine learning models at all stages of drug discovery and development illustrating some of the key areas that could be modeled.

A drug discovery and development dashboard for E2E machine learning provides the go-no-go decisions based on inputs of machine learning algorithms (SVM – support vector machine; DL – deep learning; BNB – Naïve Bayesian; KNN – K-nearest neighbors; RF – random forest; ADA-AdaBoost) or a consensus.

Figure 2. Demonstrating iterative drug discovery using machine learning.

A. The prospective machine learning approach. B. Demonstration of linkage between disease, target and machine learning model using Pitt Hopkins Syndrome as an example⁹⁵.

Table 1.

Illustrating E2E machine learning: Areas of relevance to drug discovery and development with substantial data available where machine learning models have been applied.

End points modeled
Target discovery ^{14, 15}
Molecule Synthesis ^{36, 37}
Small molecule physicochemical properties ⁸⁰
Solubility ⁸¹
Drug Induced Liver Injury ⁸²
hERG ⁸³
ADME properties ⁸²
Blood Brain Barrier penetration ⁸⁴
Skin Permeability ⁸⁵
Transporters ⁴⁵
Mutagenicity ⁸⁶
Drug Induced Liver Injury ⁸²
In vivo pharmacokinetics ⁸⁷
Reproductive toxicology ⁸⁸
Formulation ⁸⁹
Environmental impact ⁹⁰
Pharmacoeconomics / cost effectiveness analysis / policy decisions ⁹¹
Clinical trial: recruiting, design, optimization, success and failure ⁶
Manufacturing ⁹²
Counterfeit drug detection ⁹³
Post marketing surveillance adverse event prediction ⁹⁴
Electronic Health Records ⁴⁰