Active-learning strategies in computer-assisted drug discovery

Daniel Reker and Gisbert Schneider

Swiss Federal Institute of Technology (ETH), Department of Chemistry and Applied Biosciences, Vladimir-Prelog-Weg 4, 8093 Zürich, Switzerland

High-throughput compound screening is time and resource consuming, and considerable effort is invested into screening compound libraries, profiling, and selecting the most promising candidates for further testing. Active-learning methods assist the selection process by focusing on areas of chemical space that have the greatest chance of success while considering structural novelty. The core feature of these algorithms is their ability to adapt the structure–activity landscapes through feedback. Instead of full-deck screening, only focused subsets of compounds are tested, and the experimental readout is used to refine molecule selection for subsequent screening cycles. Once implemented, these techniques have the potential to reduce costs and save precious materials. Here, we provide a comprehensive overview of the various computational active-learning approaches and outline their potential for drug discovery.

Introduction

The concept of iterative molecular design, synthesis, and testing forms a central pillar of drug discovery; it provides the basis for our understanding of the underlying structure–activity relation (SAR). Iterative synthesize-and-test cycles with SAR model adaptation to newly obtained activity data improve the overall quality of the designer compounds and help reduce experimentation costs. Similarly, the screening of existing compounds profits from such feedback-driven picking: within a fixed budget, adaptive screening rounds through multiple acquisition-and-test cycles can lead to significantly better solutions compared with a single large screen [1,2]. The crucial step in each learning cycle is the formulation of a well-motivated hypothesis for compound generation (*de novo* design) or compound picking (when screening from a compound pool) based on the available SAR data. The selected molecules can either be hypothesized actives or readily available compounds that will improve the model by elucidating poorly understood parts of the SAR. Commonly, an interdisciplinary team of scientists generates the new hypothesis by inferring from their expertise and medicinal chemistry 'intuition'. Therefore, any design hypothesis is easily biased towards preferred chemistry [3,4] or predisposed model interpretation [5,6]. Although expert knowledge is indisputably important for successfully guiding drug discovery projects, an unbiased perspective during the compound selection process can lead to structurally surprising chemical agents with the desired novelty, bioactivity, and physicochemical properties [7]. Moreover, with the recent advances of microfluidics-assisted integrated medicinal chemistry platforms (e.g., lab-on-a-chip systems [8]), the generation of an accurate and suitable molecular design hypothesis and, consequently, the selection of new compounds for synthesis and testing, becomes the bottleneck in an otherwise automatable optimization process [9].

Computational models act as rapid and objective decision makers in this decisive selection step (Fig. 1a) [10,11]. Active learning (also known as 'selective sampling') is an umbrella term from the field of machine learning for methods that select data points for testing and feeding back into the model [12,13]. Approximately 15 years ago, the term was introduced to drug discovery [14]. Recently, the topic has gained momentum, driven by technological advancements in small-scale organic synthesis systems and the accuracy of machine-learning prediction models. Here, we provide a comprehensive overview of investigations that have applied active-learning techniques to drug discovery. We focus on methods for finding novel chemical structures and discuss possible future directions of algorithm development and how these might help solve current challenges in computer-assisted drug design.

Corresponding author:. Schneider, G. (gisbert.schneider@pharma.ethz.ch)

http://dx.doi.org/10.1016/j.drudis.2014.12.004

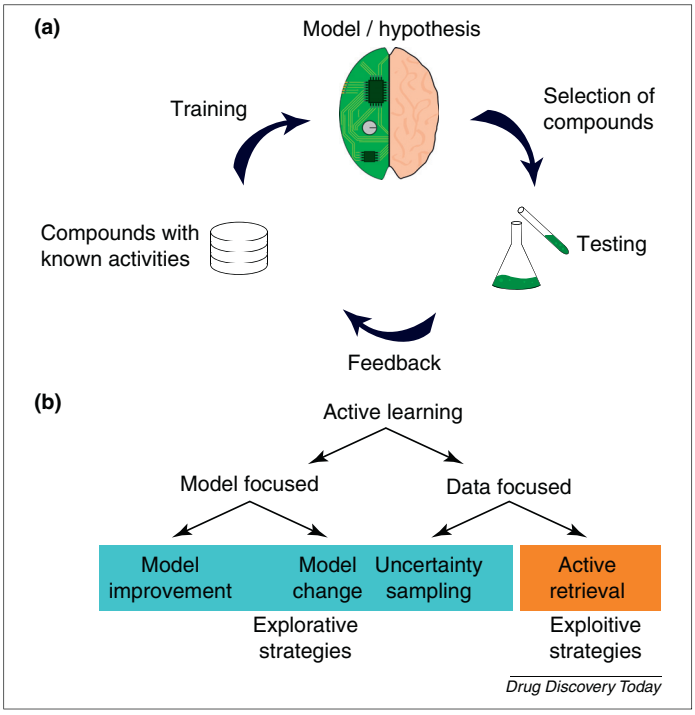

FIG. 1

(a) Schematic of the active-learning concept. Known activity data are provided as training data to a machine-learning model that generalizes this knowledge. A selection strategy is used that picks from a list of new molecules with unknown activity. These selection strategies usually try to identify molecules that would be particularly suited for improving the model quality (explorative strategies) if they are included in the training database with their activity value. Alternatively, molecules are selected that might have favorable activity values (exploitive strategies). After the selected molecules have been tested ('labeled'), they are added to the training data to train an improved machine-learning model. **(b)** Conceptual comparison of different active-learning strategies. These can be distinguished methodologically according to whether the selection strategy is derived from the whole model ('Model focused') or by examining individual data points ('Data focused'). When compounds are selected with the whole model in mind, the strategies are explorative. Possible implementations are predicting or calculating the change in model architecture ('Model change') or the improvement of the model ('Model improvement'; e.g., variance reduction or error on the test set). When examining individual data points, models can either be exploitive ('Active retrieval') or use the error or uncertainty on the individual data points to perform confined model optimization ('Uncertainty sampling').

BOX 1

Pseudocode for performing a retrospective active-learning investigation ('ActiveLearning')

```
function ActiveLearning(M,s):
    T, L, E ← splitStratified(M)
    m ← trainRFmodel(T)
    for I ← 1 to 100 do:
        selected_mol ← s(L,m)
        L ← L \ {selected_mol}
        T ← T ∪ {selected_mol}
        m ← trainRFmodel(T)
        evaluate(m,E)
    end do

function random(L,m):
    return pickOneRandom(L)

function exploitive(L,m):
    predictions ← m.predictActivitiy(L)
    return L[argmax(predictions)]

function explorative(L,m):
    uncertainty ← m.predictionVariance(L)
    return L[argmax(uncertainty)]
```

The function takes descriptions and activities of a set of molecules ('M') and a selection function ('s') that is used for the picking of molecules. First, the molecular data are split into three subsets in a stratified manner according to activity. Afterwards, the training data ('T') are used for initial model training ('trainRFmodel'). The active learning is performed for 100 iterations in which we first pick a molecule from the learning data ('L') according to the selection function. This selected molecule is then removed from the learning data and added to the training data, with which the model is retrained. The performance of the new model can then be evaluated ('evaluate'), for example according to the error on the test data ('E', Fig. 2a, main text), the activity of the picked molecule (Fig. 2b, main text), or the number of scaffolds known to the model (Fig. 2c, main text). As examples of selection functions, we show pseudocode for a random strategy ('random') that picks a random molecule from the set, an exploitive strategy ('exploitive') that picks the molecule with the highest predicted activity, and an explorative strategy ('explorative') that picks the molecule with the highest prediction uncertainty (e.g., the maximum variance according to the individual activity predictions of the trees of the random forest model).

Exploration versus exploitation

Compound selection strategies can be distinguished according to their underlying motivation (Fig. 1b): whereas some algorithms utilize the available information to retrieve compounds with certain properties ('exploitation'), others seek to improve the model by adding knowledge ('exploration'). From a technical point of view, exploration can either be performed from a molecule-centric perspective ('uncertainty sampling', i.e., selecting molecules that are predicted with low confidence by the model) or by explicitly estimating the impact of adding the additional data point on the error or architecture of the model ('model-centric' approaches). Explorative strategies sample more diverse chemical structures and rapidly increase the knowledge for the model (Fig. 2a), while not always proposing favorable structures in terms of their activity (Fig. 2b). Conversely, exploitive strategies retrieve active compounds with a greater probability, but do not necessarily add knowledge to the model. In fact, the model quality can even decrease over time when using an exploitive strategy because of the introduction of a strict bias towards highly active compounds (Fig. 2a). Various strategies for either of the two compound selection principles have been proposed and validated in the context of drug discovery (Table 1).

Explorative approaches have proven particularly attractive when aiming at novel chemotypes with desired bioactivities (Fig. 2c). For example, to probe for the applicability of uncertainty sampling to explorative drug design, Lemmen and coworkers developed both a jury of Perceptrons and a support vector machine (SVM) model to distinguish thrombin ligands from 'inactives' [14,15]. Model optimization was conducted by adding examples

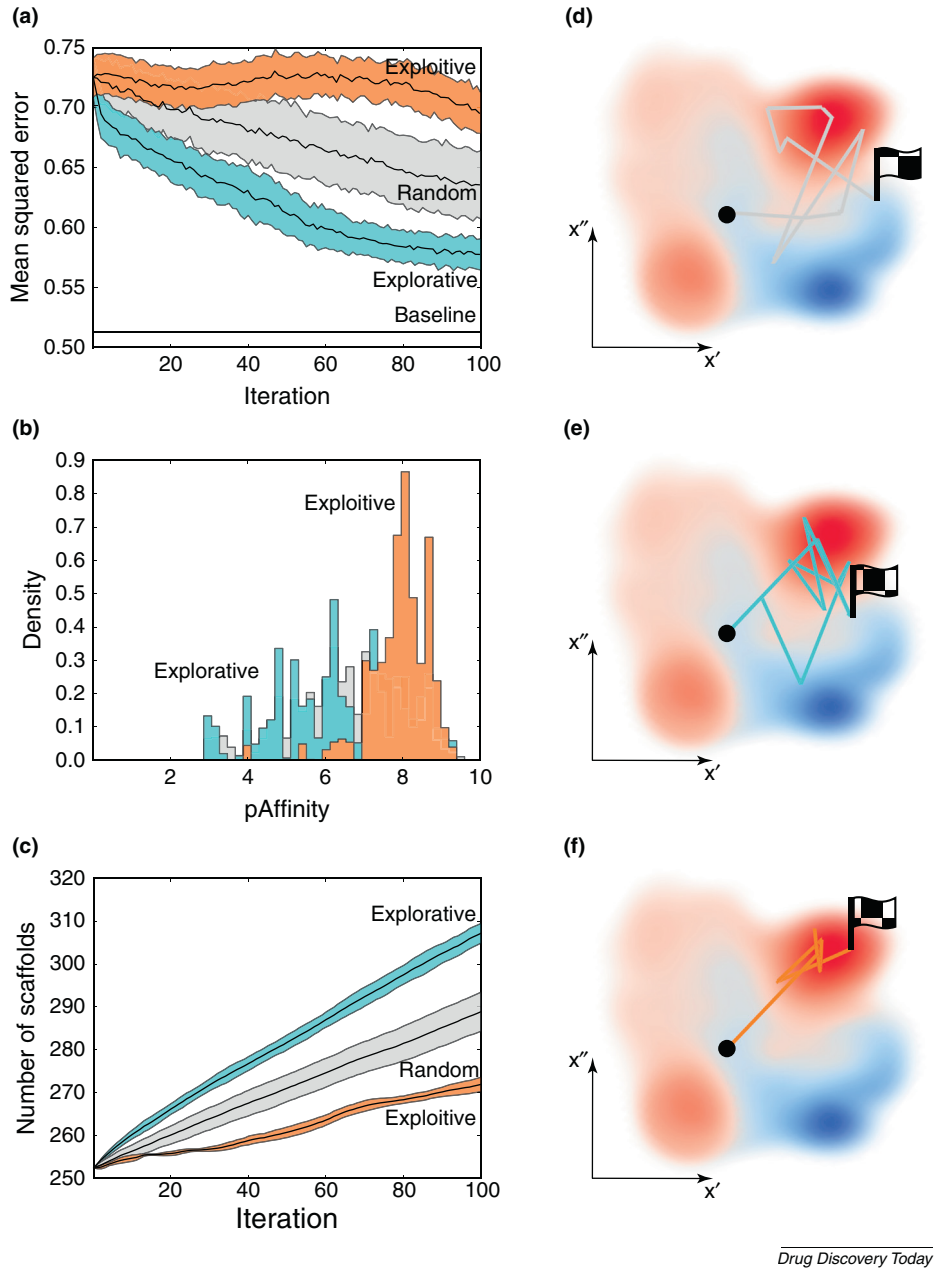

FIG. 2

Comparison of the behavior of an explorative and an exploitive strategy (see also Box 1). A random forest regression model (scikit-learn-0.14.1) was built on affinity data (IC_{50}, K_d, and EC_{50}) and Morgan fingerprints (*radius* = 4, 2048 bits) for human cyclin-dependent kinase 2 (CDK2) ligand data from ChEMBL (version 19, www.ebi.ac.uk/chembl/) [61] containing 3780 structures. The data were split into three equal parts in a stratified manner according to activity. One part was used for initial model training ('training set'), one as a set from which the model was allowed to pick new structures ('learning set') and another for external validation to monitor the development of the mean squared error (MSE) on unknown data ('test set'). Active learning was performed for a total of 100 iterations for all applied strategies. Initial model training and active learning was repeated 100 times for estimating the impact of the stochastic model creation [we show mean values and standard deviations in **(a–c)**]. Maximum prediction was used as the exploitive strategy (shown in blue) and query-by-committee (i.e., maximum prediction variance) as the explorative strategy (shown in orange). Random molecule picking (shown in gray) served as a baseline. The explorative strategy rapidly reduces the error on the test set and converges towards the minimal possible error (black line, average error of 100 models trained on all training and learning data), whereas the performance of the exploitive model fluctuates and is outperformed even by random selection (a). Conversely, the exploitive strategy successfully retrieves highly active compounds, whereas the explorative strategy samples activity equivalent to random selection (b). This is also visible in the number of scaffolds retrieved by the different strategies: whereas the exploitive strategy largely samples from the universe of known, active scaffolds, the explorative strategy selects compounds with scaffolds that are not contained in the training data (c). For further analysis, activity landscapes (Lisard-1.2.6) were created using a principle component analysis (PCA) of CATS2 descriptions of the same ChEMBL CDK2 data. Trajectories of selected molecules are visualized when the active-learning strategies are initialized with only one randomly picked example (CHEMBL326275) as training data. Whereas both random selection **(d)** and the explorative strategy **(e)** sample from larger areas of the landscape, the exploitive strategy **(f)** is focused on an activity island after it found the first highly active compound.

TABLE 1

Overview of retrospective and prospective investigations applying the active-learning concept in computational drug discovery[a]

Target	Machine learning model	Type of study	Exploration	Exploitation	Descriptor	Batch	Refs (year)
Thrombin (and CDK2)	Jury of Perceptrons; SVM	Retrospective	Uncertainty sampling	Maximum certainty	Molecular shape features	Naïve	[14,15] (2003)
GPCRs	QBag	Retrospective (+ one prospective screen)	Uncertainty sampling	None	MDL key + physicochemical properties	Naïve	[16] (2008)
Anticancer drug screen (NCI60)	Gaussian process	Retrospective	Uncertainty sampling	Variance corrected predictions and expected improvement	OpenBabel FP2	None	[18,19] (2008)
12 human targets	Gaussian process	Retrospective	None	Expected improvement	MOE 2D descriptors	None	[20] (2013)
Narcotic analgesics	KGCB	Retrospective	None	Expected improvement	Free-Wilson model	None	[21] (2011)
GPCR polypharmacology and blood–brain barrier penetration	Bayesian model	Prospective	Sampling by genetic algorithms	Maximum prediction	ECFP6	Naïve	[32] (2012)
Abl kinase	Random forest	Prospective	Undersampled building block	Maximum prediction	ECFP6 + physicochemical properties	Naïve	[34] (2013)
Gyrase	QMOD	Retrospective	Binding mode analysis	Maximum prediction	N/A	Naïve	[35] (2012)

[a] The table reports the target activity that was learned in the study ('Target') and the used machine learning algorithm ('Machine learning Model') and molecular descriptor ('Descriptor'). Furthermore, it reports whether the investigation was of retrospective or prospective character ('Type of study') and the applied active-learning selection strategies (split into strategies for 'Exploration' and 'Exploitation'). Finally, the 'Batch' column reports whether the algorithm was allowed to be informed about one selected compound immediately (None) or had to select a set of a certain number of top scoring compounds before feedback for the whole set was given (Naïve).

with the strongest disagreement among the jury or the smallest distance to the SVM model hyperplane. Both active-learning methods performed better than passive versions, particularly when the fraction of true positives was small. The authors were also able to 'invert' their selection strategy and retrieve high-confidence samples, thereby shifting the focus from model improvement to active retrieval. Fujiwara *et al*. applied uncertainty sampling through disagreement among decision trees ('query by committee', also known as 'version space reduction') in a random forest-like virtual screening approach for different G protein-coupled receptor (GPCR) ligands [16]. They analyzed individual runs and demonstrated that the explorative strategy managed to 'hop' to different classes of actives according to structural clustering and functional groups, in contrast to naïve nearest neighbor methods.

For exploitation, even such model-free approaches have been successfully applied in sequential screening campaigns [17]. For example, adding the actives found in the previous runs as additional queries generally improved the overall hit retrieval rate in subsequent screening rounds [1]. In one of the first machine-learning model-based studies, De Grave *et al*. implemented a Gaussian process model for cancer cell growth inhibition and tested several active molecule-seeking strategies [18,19]. Exploitation strategies retrieved more active molecules than did an explorative control model. However, the observed differences between most of the analyzed exploitation strategies were statistically insignificant. Ahmadi *et al*. implemented an exploitation strategy that used the expected improvement with a probabilistic notion of the difference between the predicted activity value and its implicit variance according to the underlying Gaussian process model compared with the known best compound [20]. Similar to the De Grave study, their method performed comparably to simpler exploitation strategies for many of the test cases. Not unexpectedly, the smoothness of the SAR function [10] was a crucial feature for success for the more elaborate selection strategies. This observation suggests that active-learning approaches benefit from advanced compound selection methods, but their actual impact depends on the quality of the model and the data. In fact, classical quantitative SAR (QSAR) can outperform exploitation approaches with decreasing amounts of initially provided training data [20]. With little training data at hand, the machine-learning models struggle to generalize and the retrieval of structural analogs becomes a competitive option. Nevertheless, in cases where high-quality models can be obtained, mechanistically elaborated exploitive models retrieve desirable compounds within only a few iterations. Negoescu *et al*. recently presented such a model [21]. Their method relies on an algorithmically sophisticated knowledge-gradient approach for picking compounds that maximize the expected improvement of narcotic activity.

Different selection functions operating on the same data set result in strongly differing sets of selected compounds (Fig. 2d–f), even when the strategies follow the same underlying motivation [20]. This closely resembles the situation a computational drug designer is faced with when triaging compound databases with different molecular similarity definitions, which commonly leads to different rankings and performances [22]. Transferring this knowledge to active learning suggests that it is essential to select the best-performing function for a given project according to the acquired experience with different selection functions, for

example through retrospective evaluation (Box 1). Another lesson learned from similarity searching is that a consensus combination of different perspectives might improve the general applicability and robustness of a method [23]. For example, Baram *et al.* suggested an evolving stochastic combination of compound selection functions [24]. They ran a total of 14 retrospective analyses and concluded that their adaptive consensus approach performs at least as good as the best individual model in ten out of these cases. In their study, inferior performance correlated with poor learning-data availability, which did not give the model enough room to adapt the contributions of the individual models. Nevertheless, this outcome suggests that a consensus function outperforms individual selection mechanisms and the best combination of functions can be learned during the first active-learning iterations. Evidently, the decision to change between exploration and exploitation strategies can be automated. This tactic will trigger explorative behavior when the model needs to be extended, while otherwise retrieving highly desirable compounds. Donmez *et al.* proposed a system that changes from exploration to exploitation as soon as a certain model quality is achieved [25]. Automated switching between search methods might help avoid following a 'wrong' selection strategy in active learning and increase its practical applicability.

Seeking novelty in chemical space

Finding structurally novel molecules with designer activity remains the chief aim of drug discovery [26]. The goal is not only to identify new chemical entities *per se*, but also to obtain structures with altered physicochemical and polypharmacological properties compared with the known bioactive agents. Explorative selection is not the only viable strategy, and several other possibilities for ensuring novelty in the selected compounds have been investigated.

Instead of selecting compounds from the whole compound pool, molecules can be presampled without the necessity for full-library enumeration. For example, when reaction products can be represented by their educt combination, genetic algorithms (GAs) and other adaptive sampling methods have proven useful for proposing novel molecules [1,27], including thrombin [28,29] and matrix metalloproteinase-12 (MMP-12) inhibitors [9], as well as bioactive peptides [6,30,31]. These stochastic techniques achieve explorative behavior through high 'mutation rates' or structurally diverse parents. However, they do not represent active-learning methods *per se* because they randomly explore compounds that are similar to the parent molecules in a model-free fashion and, consequently, do not include an adaptive SAR model. Nevertheless, GAs can be used in fusion approaches to ensure sample diversity, which can benefit the active-learning algorithm. Besnard *et al.* followed such a strategy using a Bayesian activity model that they provided with molecules that evolved via chemical transformations from partially 'random' parents [32]. Additionally, GAs have also been proven effective for noncombinatorial molecular representations, such as substructure or pharmacophoric fingerprints that generate descriptors that do not necessarily correspond to real molecules, and select their nearest neighbors [33]. In an orthogonal study, Desai *et al.* demonstrated that a descriptor-based random forest active-learning model was able to steer successfully combinatorial synthesis for Abl kinase inhibition, without explicitly modeling the combinatorial character of the compounds [34].

As an alternative to relying on the degree of exploration of the selection function or using presampling, researchers have also forced novelty into the selected molecules post hoc. This can be accomplished by filtering certain substructures [32] or focusing on underinvestigated building blocks [34]. Varela *et al.* followed a more elaborate approach by statistically quantifying the novelty of compounds according to their predicted receptor-binding pose [35]. Compounds were tested in parallel to inform the model about the activity of unknown structures. The authors reported that this additional testing, which obviously sampled also weakly active molecules, enabled the implemented exploitation strategy to sample structurally more diverse molecules compared with the exploitive control model.

Future developments of active learning in drug discovery

The active-learning concept has successfully been transferred from the field of computer science to drug discovery and several extensions have been proposed to improve the practical applicability of the method [12]. In the context of drug discovery, three of these theoretical considerations appear to be most promising, namely re-labeling, cost-aware learning, and batch selection.

First, readouts from biological assays are often associated with high noise levels, particularly, when data from different assays are aggregated [36]. Active-learning approaches have emerged that can challenge the annotation of known compounds by requesting retesting. Therefore, the active-learning model can aid training data curation and rescue false negatives [9].

Second, the theoretically studied concept of cost-aware learning might make active-learning algorithms more economical [37] because compounds that are difficult to synthesize, precious, expensive, or difficult to handle might be poor choices for model refinement.

Third, most biological assays are performed in batches of compounds that are tested simultaneously. All of the reported studies concerning active learning have investigated selection functions that select one compound at a time or utilize naïve selections of the top-ranked compounds according to the scoring function. Specifically for large batch numbers, this strategy has been empirically demonstrated to decrease model convergence, because the added knowledge can be redundant [16]. Given that the evaluation of every possible subset is not feasible for combinatorial reasons, heuristic strategies have been proposed to find a group of instances with little overlap in additional knowledge [38].

For complex objective functions, semisupervised learning has proven useful [39,40] and could aid batch selection. This closely related machine-learning concept also selects compounds for model refinement, but simply assumes the predicted affinities to be true. The model is then retrained with both measured and predicted activities. This type of feedback does not provide the model with more knowledge about the activities, but about the data distribution to determine more appropriate class boundaries when the data distribution is meaningful for the given classification problem and certain smoothness criteria can be met [41]. A semisupervised approach might aid in subsequently selecting near-optimal solutions and avoid redundancy by adding assumed values for the selections before additional testing.

A current hot topic in drug design is the identification of compounds that have a desired polypharmacological profile [42]. In contrast to identifying ligands that modulate one specific target activity, known as the 'magic bullet' concept [43], polypharmacological strategies have been suggested to treat complex diseases, such as cancer [44], identifying antipathogenics that have low susceptibility to developing resistance [45], as well as avoiding off-target liabilities [46]. These approaches are frequently referred to as 'magic shotgun' or 'master key compounds' to emphasize the often promiscuous character of compounds [47]. For computational polypharmacological investigations, several individual activity models are combined for target profile prediction. This adds another level of complexity to the active-learning concept because compound selection becomes a multidimensional optimization task. Hopkins and coworkers reduced dimensionality by arithmetically averaging the different objective functions and performed optimization according to the distance to the 'optimal achievement point' [32]. Using this approach, solutions are selected that perform well on all considered objectives simultaneously. Similarly, the combination of multiple objectives as a single function by combining Gaussian mixtures was recently suggested [48,49]. Both approaches led to the discovery of novel compounds that had the desired physicochemical properties and polypharmacology profiles against a selected set of GPCRs. Intriguingly, recent investigations have suggested that it would be worthwhile to account explicitly for the multidimensionality of the objective space in the developed algorithms, for example by Pareto ranking [50]. Zuluaga *et al.* suggested an active-learning concept that actively

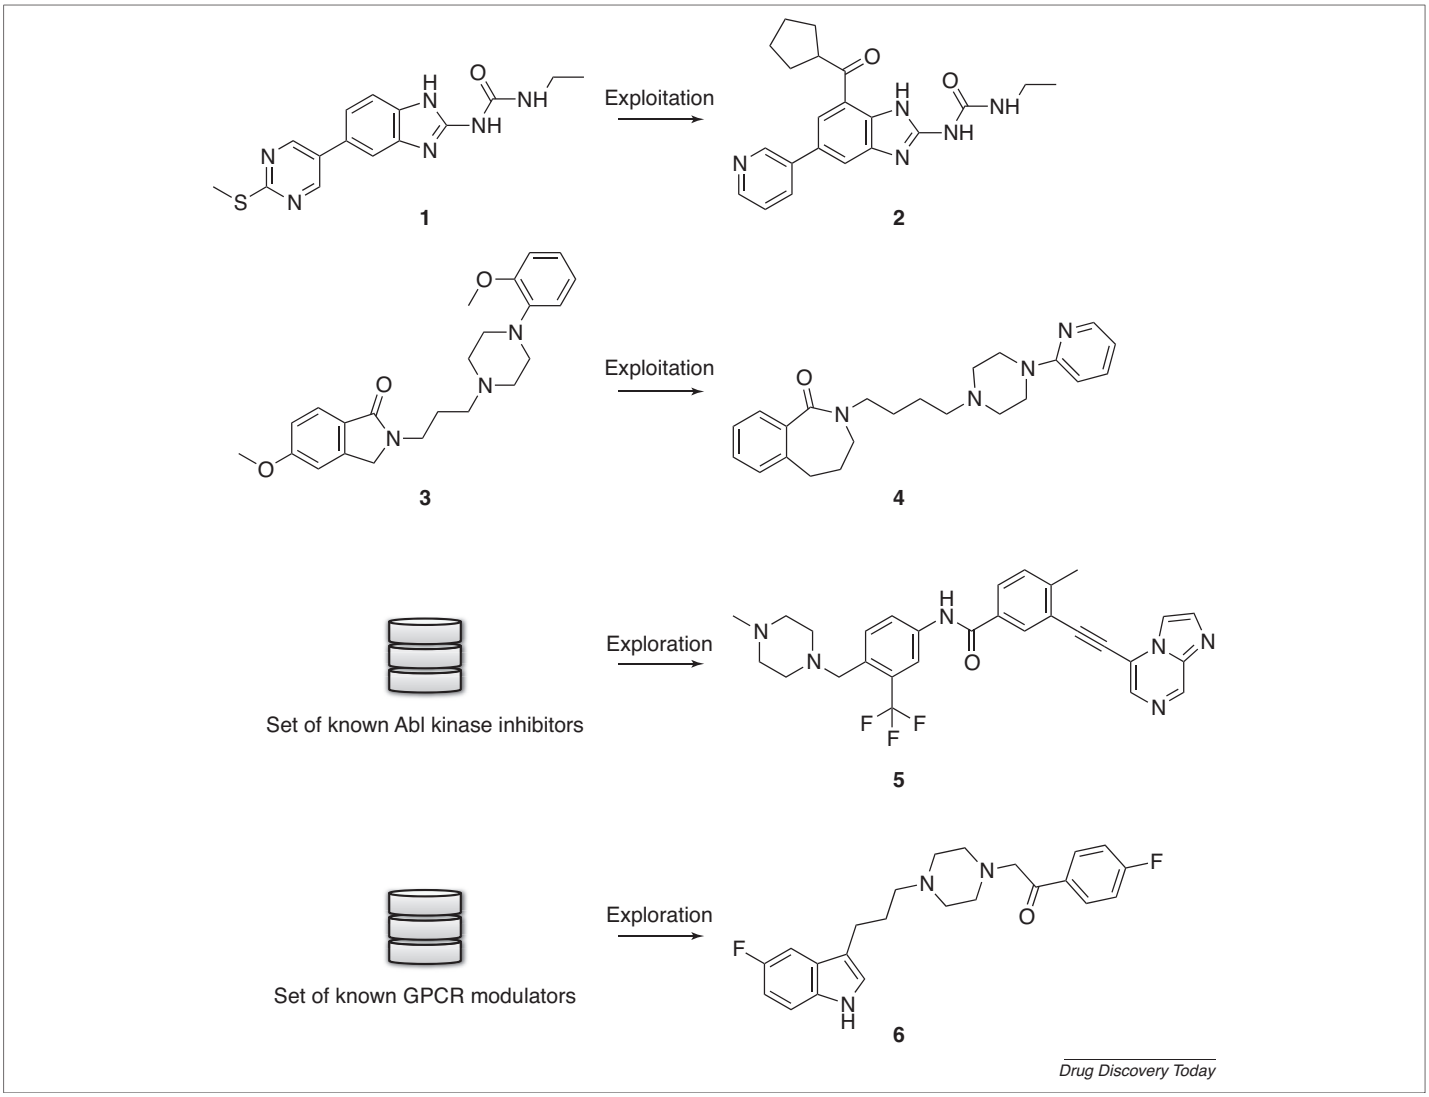

FIG. 3

Bioactive compounds identified in prospective active-learning studies, selected according to prominence and potency as reported in the respective publications. When an exploitive strategy is applied, new compounds are associated with bioactive structures from the training set in the investigated area of chemical space (e.g., **1** ← **2** [35] and **3** ← **4** [32]). Their scaffolds are largely maintained, but the molecules are unlikely to have been picked using naive similarity-based approaches. In fact, they have low pairwise Morgan fingerprint (radius = 4, 2048 bits) Tanimoto similarity (Tc; with Tc = 1 meaning identical fingerprints) to their respective reference molecule: Tc = 0.3 for **2**, Tc = 0.21 for **4**. An explorative strategy is aimed at finding novel structures (e.g., **5** [34] and **6** [16]) that are substantially different compared with the sets of known actives (maximal Tc = 0.21 for **5** against Abl kinase inhibitors, and maximal Tc = 0.14 for **6** against a set of serotonin, adrenergic, and muscarinic receptor modulators extracted from the COBRA collection [62]). Abbreviation: GPCR, G protein-coupled receptor.

strives to identify the set of Pareto-optimal solutions [51]. Using such approaches, compound selection can be steered to efficiently navigate the multidimensional objective space and assist in maturing these approaches to support polypharmacological studies.

Concluding remarks

Active learning represents a feasible and broadly applicable concept for drug discovery. Here, we have identified promising extensions for the active-learning concept, which are currently underinvestigated in practical applications. For example, batch selection is notoriously unvetted, particularly for exploitive compound prioritization [18]. Comparative studies have frequently demonstrated that simple strategies, such as nearest neighbor searching, perform well in the first iterations by retrieving closely related derivatives, but when additional new chemotypes need to be retrieved, more advanced selection methods should be pursued [14,16] (Fig. 3). Therefore, model-centric perspectives and the inclusion of cost estimates might increase scaffold diversity among the retrieved hits while actively reducing overheads. Active learning can profit from the experience gained in related fields, such as focused library design [52], diversity analysis [53], and *de novo* design [54]. Simultaneously, knowledge about elaborate active-learning strategies could be transferred from successful applications in drug formulation [55], protein–protein interactions [56], drug combinations [57], gene expression data for cancer diagnosis [58], cancer rescue mutations [59], and experiment scheduling [60]. With further algorithmic advances, this technology will soon enable automated innovative hit identification and support chemical decision making in rational drug design.

References

1 Schneider, G. and Schüller, A. (2009) Adaptive combinatorial design of focused compound libraries. *Methods Mol. Biol.* 572, 135–147
2 Schneider, G. *et al.* (2009) Voyages to the (un)known: adaptive design of bioactive compounds. *Trends Biotechnol.* 27, 18–26
3 Kutchukian, P.S. *et al.* (2012) Inside the mind of a medicinal chemist: the role of human bias in compound prioritization during drug discovery. *PLoS ONE* 7, e48476
4 Schnecke, V. and Boström, J. (2006) Computational chemistry-driven decision making in lead generation. *Drug Discov. Today* 11, 43–50
5 Claus, B.L. and Underwood, D.J. (2002) Discovery informatics: its evolving role in drug discovery. *Drug Discov. Today* 7, 957–966
6 Singh, J. *et al.* (1996) Application of genetic algorithms to combinatorial synthesis: a computational approach to lead identification and lead optimization. *J. Am. Chem. Soc.* 118, 1669–1676
7 Willet, P. (2000) Chemoinformatics – similarity and diversity in chemical libraries. *Curr. Opin. Biotechnol.* 11, 85–88
8 Rodrigues, T. *et al.* (2014) Accessing new chemical entities through microfluidic technology. *Angew. Chem. Int. Ed.* 53, 5750–5758
9 Pickett, S.D. *et al.* (2011) Automated lead optimization of MMP-12 inhibitors using a genetic algorithm. *ACS Med. Chem. Lett.* 2, 28–33
10 Bajorath, J. *et al.* (2009) Navigating structure–activity landscapes. *Drug Discov. Today* 14, 698–705
11 Hopkins, A.L. and Polinsky, A. (2006) Knowledge and intelligence in drug design. *Annu. Rep. Med. Chem.* 41, 425–437
12 Settles, B. (2010) *Active Learning Literature Survey*. University of Wisconsin
13 Guyon, I. *et al.* (2012) *Active Learning Challenge: Challenges in Machine Learning*. Microtome Publishing
14 Warmuth, M.K. *et al.* (2001) Active learning in the drug discovery process. *NIPS* 1449–1456
15 Warmuth, M.K. *et al.* (2003) Active learning with support vector machines in the drug discovery process. *J. Chem. Inf. Comput. Sci.* 43, 667–673
16 Fujiwara, Y. *et al.* (2008) Virtual screening system for finding structurally diverse hits by active learning. *J. Chem. Inf. Model.* 48, 930–940
17 Bajorath, J. (2002) Integration of virtual and high-throughput screening. *Nat. Rev. Drug Discov.* 1, 882–894
18 De Grave, K. *et al.* (2008) Active learning for high throughput screening. *Discovery Science* pp. 185–196, Springer
19 De Grave, K. *et al.* (2008) Active learning for primary drug screening. *Benelearn 08, The Annual Belgian-Dutch Machine Learning Conference* pp. 55–56; Conference Proceedings; available online at http://www.montefiore.ulg.ac.be/services/stochastic/benelearn08/
20 Ahmadi, M. *et al.* (2013) Predicting potent compounds via model-based global optimization. *J. Chem. Inf. Model.* 53, 553–559
21 Negoescu, D.M. *et al.* (2011) The knowledge-gradient algorithm for sequencing experiments in drug discovery. *NFORMS J. Comput.* 23, 346–363
22 Sheridan, R.P. *et al.* (2002) Why do we need so many chemical similarity search methods? *Drug Discov Today* 7, 903–911
23 Reker, D. *et al.* (2014) 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
24 Baram, Y. *et al.* (2004) Online choice of active learning algorithms. *J. Mach. Learn. Res.* 5, 255–291
25 Donmez, P. *et al.* (2007) Dual strategy active learning. *Machine Learning: ECML 2007* pp. 116–127, Springer
26 Medina-Franco, J.L. *et al.* (2014) Balancing novelty with confined chemical space in modern drug discovery. *Expert Opin. Drug Discov.* 9, 151–156
27 Weber, L. (1998) Evolutionary combinatorial chemistry: application of genetic algorithms. *Drug Discov. Today* 3, 379–385
28 Weber, L. *et al.* (1995) Optimization of the biological activity of combinatorial compound libraries by a genetic algorithm. *Angew. Chem. Int. Ed.* 34, 2280–2282
29 Illgen, K. *et al.* (2000) Simulated molecular evolution in a full combinatorial library. *Chem. Biol.* 7, 433–441
30 Yokobayashi, Y. *et al.* (1996) Directed evolution of trypsin inhibiting peptides using a genetic algorithm. *J. Chem. Soc., Perkin Trans. 1* 20, 2435–2439
31 Kampenhausen, S. *et al.* (2002) A genetic algorithm for the design of molecules with desired properties. *J. Comput. Aided Mol. Des.* 16, 551–567
32 Besnard, J. *et al.* (2012) Automated design of ligands to polypharmacology profiles. *Nature* 492, 215–220
33 Gobbi, A. and Poppinger, D. (1998) Genetic optimization of combinatorial libraries. *Biotechnol. Bioeng.* 61, 47–54
34 Desai, B. *et al.* (2013) Rapid discovery of a novel series of Abl kinase inhibitors by application of an integrated microfluidic synthesis and screening platform. *J. Med. Chem.* 56, 3033–3047
35 Varela, R. *et al.* (2012) Iterative refinement of a binding pocket model: active computational steering of lead optimization. *J. Med. Chem.* 55, 8926–8942
36 Kalliokoski, T. *et al.* (2013) Comparability of mixed IC_{50} data – a statistical analysis. *PLOS ONE* 8, e61007
37 King, R.D. *et al.* (2004) Functional genomic hypothesis generation and experimentation by a robot scientist. *Nature* 427, 247–252
38 Hoi, S.C.H. *et al.* (2006) Batch mode active learning and its application to medical image classification. In *Proceedings of the International Conference on Machine Learning (ICML)* pp. 417–424, ACM Press
39 Byvatov, E. *et al.* (2005) From virtual to real screening for D3 dopamine receptor ligands. *ChemBioChem* 6, 997–999
40 Franke, L. *et al.* (2005) Extraction and visualization of potential pharmacophore points using support vector machines: application to ligand-based virtual screening for COX-2 inhibitors. *J. Med. Chem.* 48, 6997–7004
41 Chapelle, O. *et al.* (2006) When can semi-supervised learning work? *Semi-supervised Learning* pp. 2–3, MIT Press
42 Paolini, G.V. *et al.* (2006) Global mapping of pharmacological space. *Nature* 24, 805–815
43 Franzt, S. (2005) Drug discovery: playing dirty. *Nature* 437, 942–943
44 Knight, Z.A. *et al.* (2010) Targeting the cancer kinome through polypharmacology. *Nat. Rev. Cancer* 10, 130–137
45 Hopkins, A.L. *et al.* (2011) Rapid analysis of pharmacology for infectious diseases. *Curr. Top. Med. Chem.* 11, 1292–1300
46 Merino, A. (2010) Drug profiling: knowing where it hits. *Drug Discov. Today* 15, 749–756

47 Medina-Franco, J.L. *et al.* (2013) Shifting from the single to the multitarget paradigm in drug discovery. *Drug Discov. Today* 18, 495–501
48 Reutlinger, M. *et al.* (2011) Neighborhood-preserving visualization of adaptive structure-activity landscapes: application to drug discovery. *Angew. Chem. Int. Ed.* 50, 11633–11636
49 Reutlinger, M. *et al.* (2014) Multi-objective molecular de novo design by adaptive fragment prioritization. *Angew. Chem. Int. Ed.* 53, 4244–4248
50 Nicolaou, C.A. and Brown, N. (2013) Multi-objective optimization methods in drug design. *Drug Discov. Today Technol.* 10, e427–e435
51 Zuluaga, M. *et al.* (2013) Active learning for multi-objective optimization. *JMLR W&CP* 28, 462–470
52 Beno, B.R. (2001) The design of combinatorial libraries using properties and 3D pharmacophore fingerprints. *Drug Discov. Today* 6, 251–258
53 Meinl, T. *et al.* (2011) Maximum-score diversity selection for early drug discovery. *J. Chem. Inf. Model.* 51, 237–247
54 Schneider, G. (2013) *De Novo Molecular Design.* Wiley-VCH
55 Caschera, F. *et al.* (2010) Automated discovery of novel drug formulations using predictive iterated high throughput experimentation. *PLoS ONE* 5, e8546
56 Mohamed, T.P. *et al.* (2010) Active learning for human protein–protein interaction prediction. *BMC Bioinform.* 11, S57
57 Small, B.G. *et al.* (2011) Efficient discovery of anti-inflammatory small-molecule combinations using evolutionary computing. *Nat. Chem. Biol.* 7, 902–908
58 Lui, Y. (2004) Active learning with support vector machine applied to gene expression data for cancer classification. *J. Chem. Inf. Comput. Sci.* 44, 1936–1941
59 Danziger, S.A. *et al.* (2009) Predicting positive p53 cancer rescue regions using Most Informative Positive (MIP) active learning. *PLoS Comput. Biol.* 5, e1000498
60 Naik, A.W. *et al.* (2013) Efficient modeling and active learning discovery of biological responses. *PLOS ONE* 8, e83996
61 Gaulton, A. *et al.* (2012) ChEMBL: a large-scale bioactivity database for drug discovery. *Nucleic Acids Res.* 40, D1100–D1107
62 Schneider, P. and Schneider, G. (2003) Collection of bioactive reference compounds for focused library design. *QSAR Comb. Sci.* 22, 713–718