

A New Generation of Possibilities: Applying med chem transformations to guide the search for high quality leads and candidates

Matthew Segall*, Ed Champness*, Chris Leeding*, Ryan Lilien†, Ramgopal Mettu†, Brian Stevens†

* Optibrium Ltd., 7226 IQ Cambridge, Beach Drive, Cambridge, UK

† Cadre Research Labs, Acton, MA, USA

Abstract

Using *in silico* predictive models and multi-parameter optimisation techniques allows large numbers of compounds to be quickly assessed with respect to a profile of properties required for a successful compound in a drug discovery project. With these predictive methods, it becomes possible to consider a large number of ideas for potential compounds that can be easily created and entered into a computer by an individual. In this article we describe a method that automatically generates chemically relevant compound ideas from an initial molecule, based on medicinal chemistry 'transformation rules' taken from examples in the literature. These are then prioritised using *in silico* models and a probabilistic scoring algorithm to identify the compound ideas most likely to satisfy a user-defined profile of required properties. Embedded in an intuitive, visual user interface, this approach provides a powerful means to explore potential chemistry to identify high quality leads or to improve properties in lead optimisation. We demonstrate that the set of 206 transformations employed is generally applicable, produces a wide range of new compounds and is representative of the types of modifications previously made to move from lead-like to drug-like compounds. Furthermore, we show that more than 94% of the compounds generated by transformation of typical drug-like molecules are acceptable to experienced medicinal chemists. Finally, we illustrate an application of our approach to the lead that ultimately led to the discovery of Duloxetine, a marketed serotonin reuptake inhibitor. Our analysis results in the identification of a diverse range of high scoring compounds, including Duloxetine itself.

Introduction

In silico predictive models of key properties are routinely used in the selection and design of potential drug molecules (1). These results may be combined to prioritise compound ideas for synthesis, simultaneously optimising multiple parameters to identify compounds with an appropriate balance of properties for the therapeutic goal of a drug discovery project (2)(3). Furthermore, the structure-activity relationships that these models capture can guide the redesign of compounds to improve their properties and overcome liabilities (4).

Predictive methods can score and rank compounds to guide the search for high quality compounds among a large number of possibilities; therefore, getting the maximum value depends on having a rich set of potential compounds to search. However, during optimisation it is rare for a large library of relevant, predefined structures to be available and it is common to rely on a medicinal chemist to define possible compounds of interest, either by drawing individual structures or enumerating

virtual libraries based on a common structural motif. This is a time consuming process and limited by the experience of an individual chemist; how many ideas can one person generate?

Methods for automatically applying medicinal chemistry 'transformation rules' to generate new compound structures have been previously described (5)(6). These typically accept an initial 'parent' structure as input and generate 'child' structures by applying transformations based on collective medicinal chemistry experience. Examples of transformation rules range from simple substitutions or bioisostere replacements to more dramatic modifications of the molecular framework such as ring opening or closing. A computer can store and apply many more rules than a single chemist and can 'learn' from historical examples of transformations between molecules (7). Applying a set of transformations iteratively to produce multiple 'generations' of compound ideas can result in a large number of molecules – too many to be examined visually by a chemist to select the most interesting for further consideration.

In this paper, we describe the combination of an algorithm to generate compound ideas, by applying transformations to an initial molecule, with predictive models and a multi-parameter scoring algorithm to quickly focus attention on those ideas most likely to satisfy the required property profile. The goal is a tool to support experts and stimulate the process of innovation – achieving a creative combination of a computer's ability to cover a wide breadth of possibilities with the experience and detailed knowledge of a chemist. In particular, the discovery process should be directed by an expert and provide a prioritised list of possibilities for further consideration, not an automatically designed final compound.

To be successful, such a method must satisfy a number of requirements:

- It must generate a wide diversity of chemistry, as the objective is to explore many ideas in the search for an optimal solution.
- The compound structures generated must be relevant. In particular, the number of 'nonsensical', e.g. chemically unstable or infeasible, compounds must be kept to a minimum. Also, the chemist must be able to control the generation process, for example by specifying a region that must not be modified or restricting the transformations that will be applied.
- The transformations that are applied should include a broadly representative set of those applied successfully in the past to optimise successful drugs.
- The method used to prioritise the resulting compound ideas should reliably identify high quality compounds within those given the highest rank in the generated set.

In the following sections, we will describe the methods used to create and apply a set of transformations and prioritise the compounds generated thereby. Furthermore, we will describe the validation of this method to ensure that the transformations cover a broad range of 'drug like' chemistry and that the resulting structures are relevant and not unstable or infeasible. Finally, we will describe the application of our method to efficiently identify compounds similar to known drugs, starting from the lead compounds from which the drugs were derived. Although retrospective, this application will demonstrate the ability to efficiently target high quality compounds.

Methods

Transformations

Two hundred and six transformations were generated manually, as SMIRKS codes, by study of medicinal chemistry literature (8)(9)(10)(11)(12)(13)(14)(15)(16)(17)(18)(19)(20)(21)(22)(23)(24) and observation of the optimisation steps between known drugs and the lead molecules from which they were derived. SMIRKS is a reaction transform language designed by Daylight Chemical Information Systems which uses SMILES and SMARTS notations to specify a generic reaction or transformation (25).

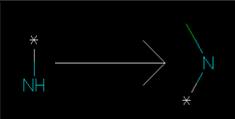
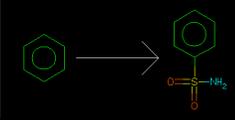
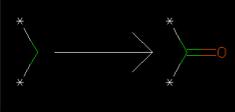
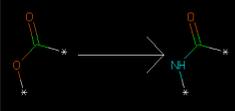
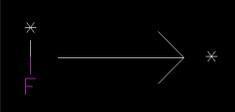
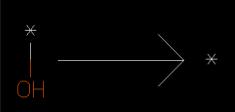
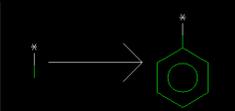
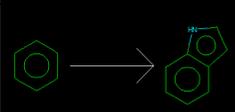
The transformations were divided into seven broad groups: Functional Group Addition, Linker Modification, Remove Atom, Ring Addition, Ring Modification, Ring Removal, Terminal Group Exchange. The distribution of transformations between the groups is shown in Table 1 and examples of each are shown in Table 2.

Table 1 Distribution of transformation between groups.

Group	Number of transformations
Functional Group Addition	20
Linker Modification	54
Remove Atom	5
Ring Addition	13
Ring Modification	26
Ring Removal	4
Terminal Group Exchange	84
Total	206

The transformations do not necessarily correspond to specific chemical reactions or synthetic routes; rather they are intended to describe changes to molecules that a medicinal chemist might consider in the course of an optimisation project. A single transformation might require multiple synthetic steps or the synthesis of new building blocks. However, the transformations are typically not major rearrangements – they are relatively feasible moves in chemical space.

Table 2 Example Transformation Rules.

Group	Transformation Name	Illustration	SMIRKS
Functional Group Addition	Methyl addition to amine		<chem>[N:1][H]>>[N:1]C</chem>
	Sulfonamide addition to benzene		<chem>[c:1]1[c:2][c:3][c:4][c:5][c:6]1[H]>>[c:1]1[c:2][c:3][c:4][c:5][c:6]1S(N)(=O)=O</chem>
Linker Modification	Secondary carbon to carbonyl		<chem>[*;!#1:1][CH2][*;!#1:2]>>[*;!#1:1]C(=O)[*;!#1:2]</chem>
	Ester to amide linker		<chem>[#6:1]O[C;!R:3](=O)[#6:2]>>[#6:1]N[C;!R:3](=O)[#6:2]</chem>
Remove Atom	Remove halogen		<chem>[C,c:1][F,Cl,Br,I]>>[C,c:1]</chem>
	Remove hydroxyl		<chem>[C,c:1][OH]>>[C,c:1]</chem>
Ring Addition	Methyl to phenyl		<chem>[*;!#1:1][CH3]>>[*;!#1:1]c1ccccc1</chem>
	Benzene to indole		<chem>[c:1]([H])1[c:2]([H])[a:3][a:4][a:5][a:6]1>>[C:1]12[a:6]=[a:5][a:4]=[a:3][C:2]=1[nH]C=C2</chem>

Ring Modification	Phenyl to 3-pyridine		<chem>[*;!#1:1][c:2]1[c:3][c:4][c:5][cH][c:6]1>>[*;!#1:1][c:2]1[c:3][c:4][c:5][n][c:6]1</chem>
	NC-switch		<chem>[*:1]1:[c]([*:2]):[c:10]([*:3]):[n]([*:4]):[*:5]1>>[*:1]1:[n]([*:2]):[c:10]([*:3]):[c]([*:4]</chem>
Ring Removal	Napthalene to benzene		<chem>[*;!#1:7][c:1]1[cH]c2c([cH][c:6]1)[c:5][c:4][c:3][c:2]2>>[*;!#1:7][c:1]1[c:2][c:3][c:4][c:5][c:6]1</chem>
	Remove phenyl		<chem>[*;!#1:1]c1[cH][cH][cH][cH][cH]1>>[*;!#1:1]</chem>
Terminal Group Exchange	Carboxyl to amide		<chem>[*;!#1:1][C:2](=O)[OH]>>[*;!#1:1][C:2](=O)N</chem>
	Amide to sulfonamide		<chem>C(=O)([NH2])[*;!#1:1]>>S(=O)(=O)([NH2])[*;!#1:1]</chem>

Generation of Compound Structures

The Cactvs cheminformatics library (26) was used within the StarDrop software platform (27) to apply the transformations, encoded as SMIRKS, to a parent compound structure encoded as a SMILES string. The Cactvs implementation also allows a fragment of the parent to be specified as a SMARTS pattern, such that this fragment will not be modified during the generation process and any transformations that would modify this region will be ignored.

The user can specify the parent structure and control the generation process through a graphical user interface implemented in the StarDrop software platform. The typical workflow is illustrated in Figure 1: The user can specify a region of the compound that must not be modified; the transformations to be applied can be selected; the number of generations of transformations to be applied can be specified; and finally, because this process generates a number of compounds that grows exponentially with the number of generations, the user can control this growth by specifying a property criterion to select a subset of the compounds in each generation. The criterion may be defined in terms of any predicted property or a score that represents the overall quality with respect to a profile of properties (see "Scoring" below) and can be specified as a threshold value for the property, e.g. only accept compounds with $\log S > 1$, or the number or proportion of compounds to select from a list ranked by the property, e.g. only progress the 100 most potent compounds in a generation or the highest scoring 10% of a generation.

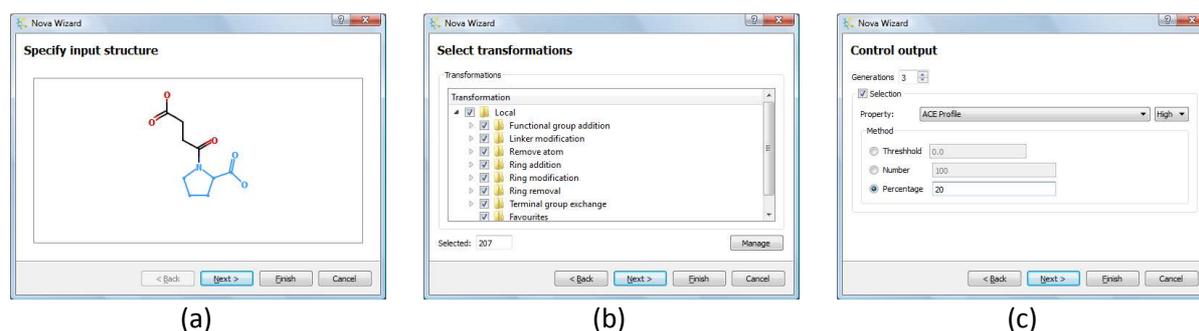


Figure 1 Illustration of workflow to initiate the generation of new compound structures. (a) Specify the input structure. A region of the molecule can be chosen to be 'frozen' (shown in light blue), in which case no modifications will be made to this region. (b) The transformations to apply can be selected, either individually or as groups. The groups can be managed to create groups tailored to specific objectives or to add new transformations. (c) The number of generations can be specified and a criterion for selection can be defined to limit the growth of the number of compounds generated. The selection can be defined as a minimum threshold for a property or score or a maximum number or percentage of each generation that will be used as the basis for subsequent generations.

Predictive Models

In principle, any *in silico* model may be used to predict the properties of the compounds generated. However, due to the large number of compounds that may be generated, the models should be capable of generating predictions quickly in order to prevent the process from becoming intractable.

In the examples presented in this paper, quantitative structure-activity relationship (QSAR) models implemented in the StarDrop software platform (27) were used to predict the following absorption, distribution, metabolism and elimination (ADME) and physicochemical properties: octanol/water partition coefficient ($\log P$), aqueous solubility ($\log S$), human intestinal absorption (HIA), blood-brain barrier penetration ($\log BB$), inhibition of the potassium ion channel encoded by the human ether-a-

go-go related gene (hERG pIC₅₀), human plasma protein binding (PPB), inhibition of cytochrome P450 isoforms CYP2D6 and CYP2C9 (pKi) and active transport by P-glycoprotein (P-gp).

In order to identify high quality compounds it is also necessary to predict activity against the pharmacological target for the intended drug. In the example application described herein, a QSAR models of target potency (expressed as the logarithm of the K_i in nM) was generated using the AutoModeller implemented in StarDrop (28) using a Gaussian Processes method (29). The data set used to build this model was derived from the ChEMBL database provided by the European Bioinformatics Institute (30). The resulting model has an R² of 0.81 and a root mean square error of 0.76 on an independent test set of 311 compounds.

Scoring

The methods underlying the probabilistic scoring algorithm employed herein are discussed in more detail in (3)(4) but here will give a brief overview. A probabilistic score is one which indicates the probability of success of a molecule against a 'scoring profile' that defines criteria for the properties that are required in an ideal compound. It is also important to specify the relative importance of the criteria as, in practice, it is often necessary to make a trade-off between properties if an ideal molecule cannot be identified. Furthermore, more subtle trade-offs can be defined than simple pass/fail criteria, as a scoring profile could contain more complex functions for each property representing a range of acceptability over the property value range. An example of such a scoring profile is shown in Figure 2.

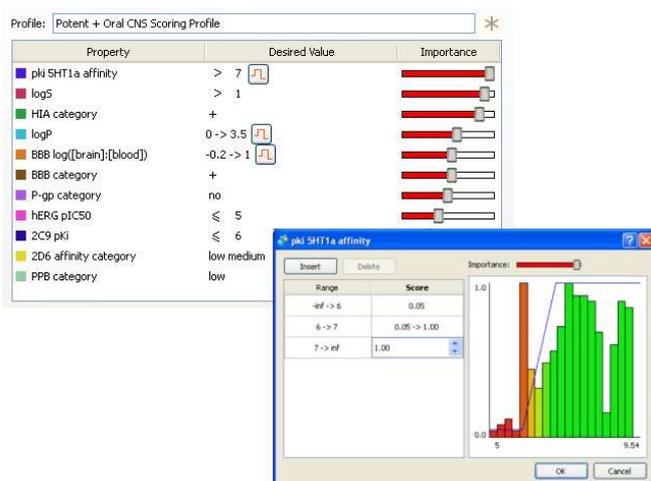


Figure 2 A 'scoring profile' showing the properties of interest, the project's success criteria and the importance of each to the project's objective. The inset window shows how more subtle trade-offs than simple pass/fail criteria can be defined, in this case a range of values over which the property value goes from ideal to unacceptable.

When combining property data on multiple properties, it is also important to consider the uncertainty in each data point, as this could lead to the overall uncertainty in the scores being high, reducing our ability to confidently distinguish high and low quality molecules. The result of this process is a score for each molecule, representing the likelihood of a molecule meeting the scoring criteria and an uncertainty in the overall score, derived from the uncertainties in each of the individual property values. These uncertainties can be used to establish whether the available data allow one molecule to be confidently chosen over another. An illustration of the output for a small set of molecules is shown in Figure 3.

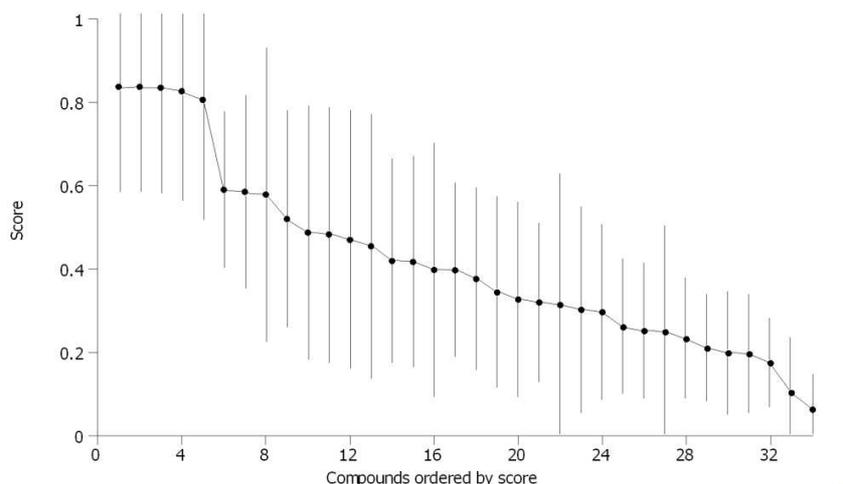


Figure 3 In this graph, molecules are plotted along the x-axis in rank order. The score is plotted on the y-axis, with error bars indicating the overall uncertainty in the score. Here the top 5 compounds cannot be confidently distinguished; more data or further criteria are required to choose between these. However, ~50% of compounds are significantly less likely to meet the project criteria than the top 5.

Visualisation

Due to the large number of compounds and volume of associated data that this process can generate, it is important to provide visual tools to guide the exploration of the rich data set generated. In addition to typical scatter plots and histograms it can be valuable to explore the parent-child relationships between generated molecules to identify transformations that have a large impact on predicted properties. An example of such a visualisation is shown in Figure 4.

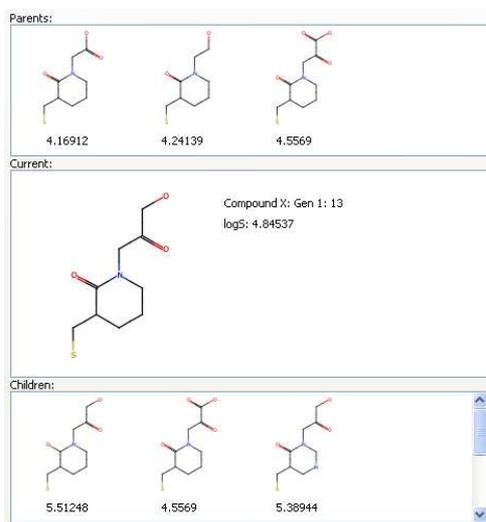


Figure 4 A view of the relationships between compounds in a dataset generated by the algorithm. The currently selected compound is shown in the middle, the parent compounds from which it was generated by different transformations are shown above and child compounds are shown below. The network of related compounds can be navigated by selecting compounds above or below the current compound. The value of a property, in this case $\log S$, is shown with each compound allowing transformations that give rise to large changes in the property to be easily identified.

It may also be useful to visualise the diversity of the compounds generated and trends in properties and scores across this diversity. An example of such a 'chemical space' is shown in Figure 5.

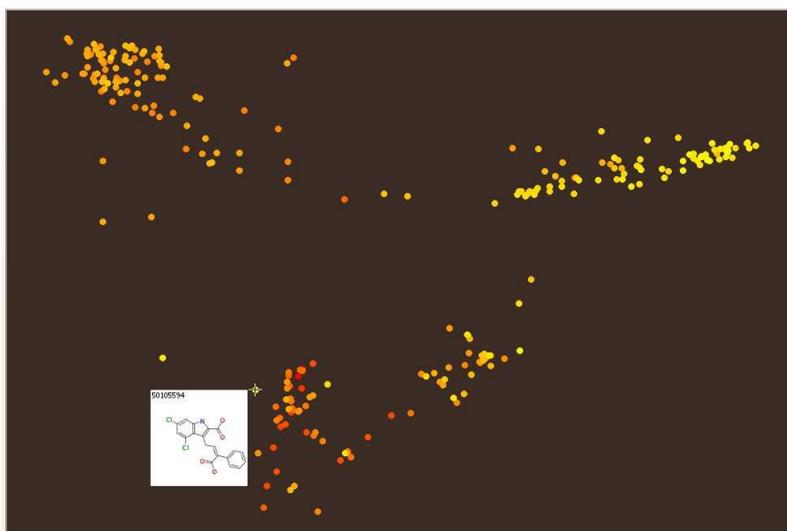


Figure 5 An example of a 'chemical space' visualisation. In this plot, each point represents a compound and the distance between two points indicates their structural similarity; close points are structurally similar while distant points are structurally diverse. The colour scale shows the distribution of a property or score. In this case the best compounds are shown in yellow and the worst in red, indicating a 'hot spot' in the top right where the best compounds are concentrated.

Transform Set Validation

Coverage

In order to ensure that the set of transformations covers a wide range of 'drug-like' chemistry, enabling the exploration of a diverse range of potential modifications, each transformation should apply to a wide range of molecules; a transformation that uniquely applies to a single molecule is not of interest. Furthermore, when the full set of transformations is applied to a 'typical' drug-like parent molecule, a large number of child molecules should be generated.

To test these requirements, the 206 transformations were applied to a set of 3,211 drug molecules (the "drug" set) derived as follows: Version 2.5 of the DrugBank Small Molecule database (31) was obtained on August 23, 2010. This initial set containing 4854 molecules was reduced by removing molecules containing atoms other than C, H, N, O, P, S, Cl, or F, molecules with molecular weight less than 200 Da and 140 molecules which contained poorly specified SMILES (127 aromaticity errors and 13 valence errors), resulting in 3214 compounds. Finally, 3 additional molecules (insulin, inulin and DB05413) were removed, as these are very large, not representative of the compounds to which we expect this method to be applied and likely to skew the validation statistics due to their size. 40 compounds were slightly edited to remove small cofactors or counter-ions or to select only one isomer where multiple isomers were specified.

The 206 transformations were applied to the drug set resulting in 584,124 child compounds; thus, on average, 182 child compounds were generated from each parent. Furthermore, on average, each transformation applied at least once to 31% of the molecules in the drug set.

These statistics indicate that the set of transformations have broad applicability to drug-like compounds and will generate a wide range of child compounds.

Quality

As discussed above, the transformation rules should be sufficiently general. However, there is a trade-off in that a more general transform is more likely to apply in an occasionally inappropriate chemical context. This can generate undesirable or infeasible compound structures. The desirability of compound structures is, to some extent, subjective. Therefore, the quality of the compound structures generated was assessed by asking two independent medicinal chemists to examine a set of 1,500 compounds generated using the 206 transformations.

The quality assessment set was generated as follows: 400 compounds were randomly selected from the drug set described above. All of the 206 transformations were applied to the 400 selected molecules to generate a set of child compounds. From the full set of child compounds, 1500 were selected at random for assessment by the medicinal chemists.

The medicinal chemists were asked to assess each child compound to determine whether it was undesirable. They were not asked to determine if they could identify a synthetic route to the product – an ideal compound that was synthetically challenging may be worth the effort of devising a difficult synthetic route or may spark further ideas that are more accessible.

From the same set of 1500 child compounds, one chemist flagged 7% of the structures as undesirable while the other flagged 4.1%. This demonstrates that desirability is, to some extent, subjective. However, an average acceptance rate of 94% was considered to be more than sufficient. It would be possible to filter out some of the undesirable structures before they are output. However, it was decided to retain this small proportion of poor compound structures as, though they may be a minor distraction, they may stimulate ideas for similar compounds that are chemically feasible.

Hit-like to Drug-like Transformation Series

The transformations in the set should be representative of those used in practice to optimise leads into drug molecules. To assess this, a data set containing 60 marketed drugs and the initial leads from which they were derived, published by Perola (32), was used (we will refer to these lead/drug pairs as the “Perola” set).

For each lead/drug pair in the Perola set, the lead was used as the initial parent and the 206 transformations were applied iteratively to explore the ‘universe’ of compounds that are accessible from the lead. The goal of this was to identify the closest compound structure in this universe to the corresponding drug. This is challenging, as many of the derivations of drugs in the Perola set from their corresponding leads include the exchange or incorporation of large or relatively uncommon fragments. A result of the coverage requirements described above is that most of the transforms involve smaller fragments. Therefore, many iterative applications of the transformations may be required, creating many generations of child compounds, to move from a lead to a compound similar to the corresponding drug and, even then, it may not be possible to find an exact match to the drug.

As the number of compounds generated increases exponentially with the number of generations, it is impractical to exhaustively enumerate all offspring compound structures. For example, if 182

compounds are generated on average from a single parent, the third generation will contain more than 6 million compounds. Therefore, a 'beam' search was implemented, whereby the 100 compounds with the greatest similarity to the target drug were retained after each iteration and a total of five iterations were applied. The closest match to the corresponding drug was identified from the resulting child compounds. The disadvantage of this approach is that it does not guarantee to find the closest match that could be achieved, as it may be necessary to initially move away from the drug in order to ultimately generate the most similar compound. Furthermore, it may be possible to find a closer child compound if more than five iterations were applied.

Similarity was measured using the Tanimoto index calculated between topological path-based fingerprints, with a maximum path length of 7 and a fingerprint size of 2048 bits. This was performed using the RDKit toolkit (33).

Out of the 60 Perola lead/drug pairs, 7 exact matches were achieved within the compounds generated from the initial lead. On average, the similarity of the drug with closest match in the child compounds generated from the corresponding lead was 0.85 compared with an average similarity between the drugs and leads of 0.64. The structures of the initial leads, corresponding drugs and closest identified child compounds are provided in the Supporting Information. This demonstrates that the transformations are representative of those used to move from lead-like to drug-like compounds.

Example Application

To illustrate the application of the transformation set to guide the search for optimised compounds based on an initial lead, we used the lead molecule that ultimately gave rise to the drug Duloxetine as the parent molecule.

The ADME QSAR models described above and a model of the inhibitory constant K_i for the serotonin transporter were used to prioritise the compounds generated against the scoring profile shown in Figure 6, which combines potency against the primary target with suitable ADME properties for an orally dosed compound against a CNS target.

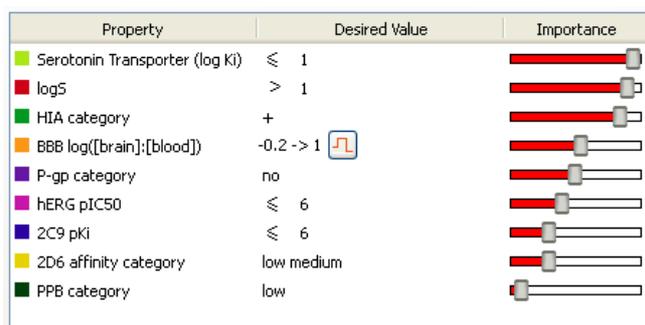


Figure 6 The scoring profile used to prioritise compounds generated from the Duloxetine lead, showing the properties of interest, the desired value ranges and the importance of each criterion. For example, the most important property was inhibition of the serotonin transporter, for which a predicted K_i of less than 10 nM ($\log K_i < 1$) was required. This was followed by an aqueous solubility of greater than 10 μM ($\log S > 1$) and positive prediction for human intestinal absorption.

The application of one generation of transformations produced 172 child compounds, which suggested that exhaustive enumeration of more than two generations would be intractable. Therefore, three generations were applied, but only the top-scoring 10% of the compounds in each of generations 1 and 2 were used as the basis for subsequent generations.

The resulting data set contained 2,208 compounds (all of the compounds in the final generation were retained) and the scores for these compounds are plotted in Figure 7. From this, a number of observations may be made: The compounds in each generation typically show an increase in score over the previous generation; the score for the initial lead is 0.09 and the averages for the compounds in subsequent generations are 0.32, 0.44 and 0.53 respectively (note that only the top 10% of the first two generations are included). However, as the results from multiple uncertain predictions are combined to calculate the score, the uncertainties in the score are high, as shown by the error bars in Figure 7. Therefore, it is difficult to discriminate between compounds with confidence, particularly in the later generations. Finally, it is notable that Duloxetine itself is present in the final generation, with a score that is significantly higher than the initial lead (level of significance ~ 0.1) and not significantly below that of the highest scoring compounds.

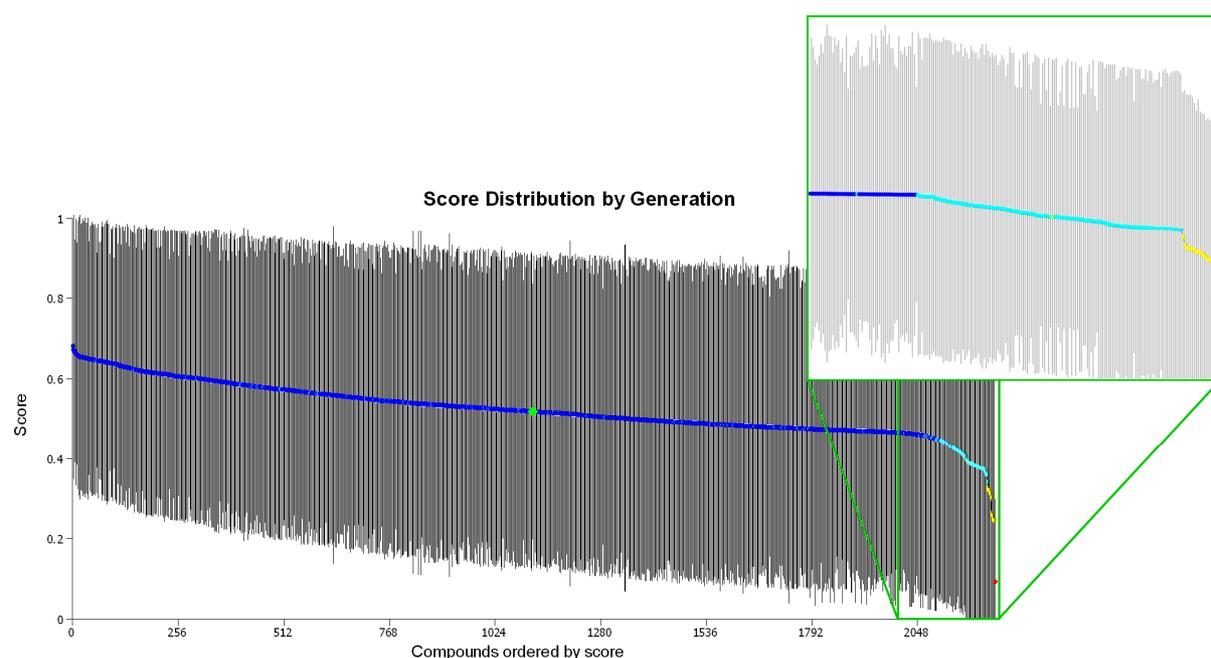


Figure 7 This graph shows the compounds generated by three generations of transformations starting with the lead compound for the project that yielded the drug Duloxetine. Error bars show the uncertainty of the overall score for each compound due to the uncertainties in the underlying data. Only the top 10% of generations 1 and 2 were used as the basis for subsequent generations. The compounds are coloured by generation: Red is the parent, yellow generation 1, light blue generation 2 and dark blue generation 3. The drug Duloxetine was present in generation 3 and is shown by the green diamond.

The structures and scores of the initial lead and Duloxetine are shown in Figure 8 along with the three highest ranking molecules generated. Although none of the top-three compounds could be identified in a search of PubChem (34), the second-ranked compound bears a strong similarity (Tanimoto similarity >0.9) to Litoxetine, shown in Figure 9, which was progressed to clinical trials and is active against the serotonin transporter with an IC₅₀ of 6 nM (35).

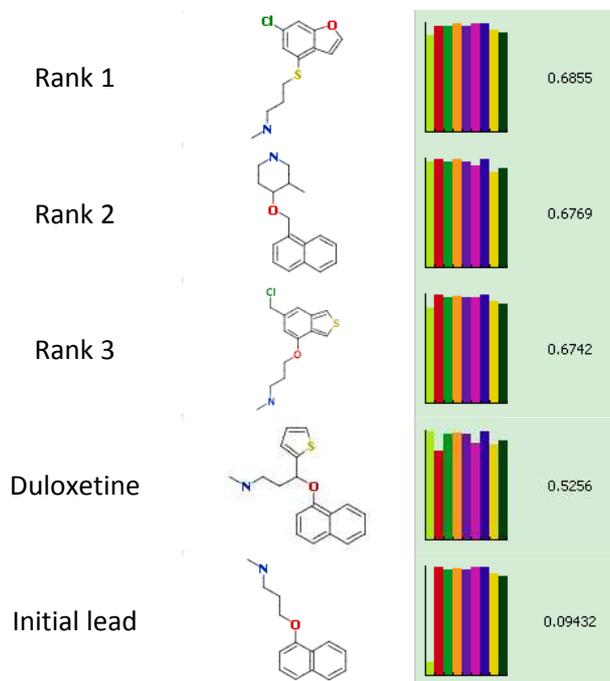


Figure 8 The initial lead that ultimately gave rise to Duloxetine, the top three compounds generated from this lead and Duloxetine, which was also generated by the algorithm. The score for each compound is shown to the right along with a histogram indicating the contribution of each property to the overall score (the colour of each bar corresponds to the property key shown in Figure 6). All of these compounds are predicted to have good values for the predicted ADME properties. However, the initial lead has a much lower score due to a significantly poorer K_i predicted for the serotonin transporter.

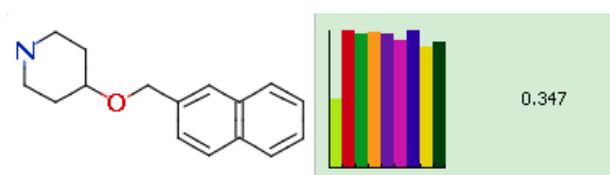


Figure 9 The structure and calculated score for Litoxetine, a clinical candidate serotonin reuptake inhibitor. The predicted K_i for this compound is 10 nM, in line with the reported IC₅₀ of 6 nM. Although this structure was not generated automatically in this example, it bears a strong similarity (Tanimoto similarity >0.9) with the second-ranked compound, which has a higher predicted affinity and hence a higher score.

The chemical space of the data set generated is shown in Figure 10. From this it is notable that a wide range of different chemical motifs have been explored and that there are multiple 'hot spots' containing high-scoring compounds; the best scoring compounds are not concentrated in one region, indicating that the algorithm has identified a number of different chemical strategies worthy of further consideration. The top three ranked molecules are structurally diverse, within the range of diversity explored around the initial lead, and are distinct from both the initial lead and Duloxetine itself.

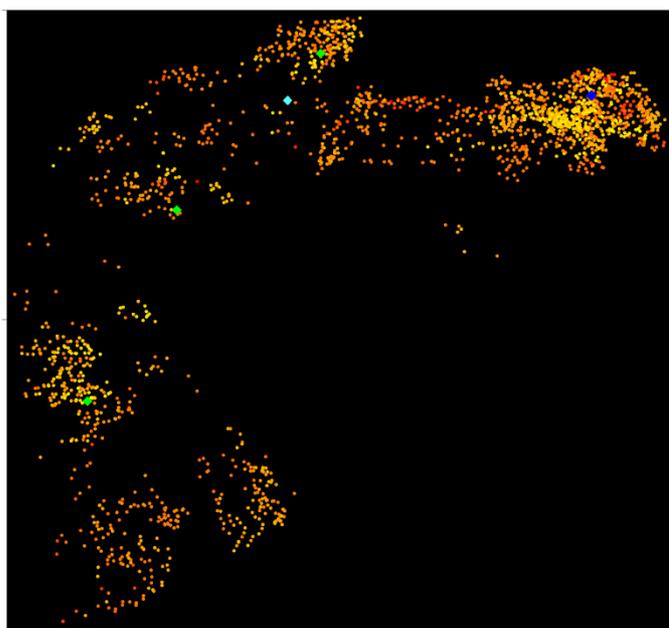


Figure 10 The chemical space of compounds generated from the initial lead that gave rise to Duloxetine. The points corresponding to compounds are coloured by score, from the lowest (0.29) in red to the highest (0.69) in yellow. The initial lead is shown as a dark blue diamond, Duloxetine as a light blue diamond. The top-three scoring compounds are shown as green diamonds.

In this example, the increase in score is driven primarily by the improvements in predicted target affinity between generations because the predicted ADME properties of the lead compound were good to begin with. However, the use of probabilistic scoring to select compounds with a good balance of properties was valuable as it eliminated compounds in early generations that were predicted to have high target affinity but were unlikely to have a good balance of ADME properties for the overall objective. Figure 11 shows the distribution of the scores for compounds in the first two generation with predicted K_i less than 10 nM, indicating that a significant number of compounds that were predicted to be active were rejected due to the predictions of poor values of other properties including solubility (184 compounds from generation 2 were used as the progenitors of generation 3).

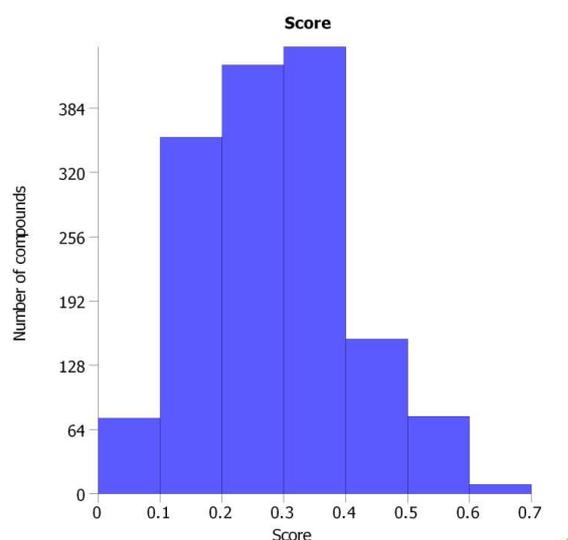


Figure 11 Score distribution for the compounds in generations 1 and 2 from the Duloxetine lead compound with a predicted K_i of less than 10 nM. From this we can see that there are a significant number of compounds with poor scores, despite having high target affinity, indicating that they are likely to have poor values for other relevant properties.

Conclusions

In this paper we have described an algorithm for automatically generating new compound ideas from an initial molecule using a set of medicinal chemistry transformations derived from the literature. We have shown that these transformations are generally applicable and generate structures that are relevant and acceptable to medicinal chemists. Furthermore, we have demonstrated the use of this chemical transformation algorithm coupled with predictive models and a multi-parameter optimisation method, integrated in an intuitive visual environment, to stimulate the exploration of a wide range of strategies to identify compounds with a good balance of properties and hence a high chance of downstream success.

There are a wide range of potential applications of this technology. These include: aiding the rigorous exploration of chemistry around early hits, to identify those hits most likely to yield high quality lead series; helping to find strategies to overcome problems with compound properties in lead optimisation; and identifying patent busting opportunities by expanding the chemistry around existing development candidates or drugs to search for compounds with improved properties.

Finally, while we have focussed on the creation and validation of an initial set of transformations, it is possible to extend this set with new transformations based on the experience of medicinal chemists or designed around specific chemistry available within an organisation. Furthermore, it may be beneficial to organise transformations into groups, perhaps tailored to specific objectives such as improving metabolic stability or reducing plasma protein binding. Thus, this approach could be used as a tool to capture and share knowledge between medicinal chemists or even as an educational resource for less experienced scientists.

Acknowledgements

The authors would like to thank Miss Delphine Lariviere for generating the model of serotonin transporter inhibition used in the example application. We would also like to thank Dr Nawaz Khan and Dr Colin Greengrass for each visually inspecting 1,500 molecules to independently assess their acceptability.

References

1. *ADMET in silico modelling: towards prediction paradise?* **Van de Waterbeemd, H and Gifford, E.** 2003, *Nat. Rev. Drug Discovery*, Vol. 2, pp. 192-204.
2. *Towards a new age of virtual ADME/TOX and multidimensional drug discovery.* **Ekins, Sean, et al.** 2002, *J. Comput. Aided Mol. Des.*, Vol. 16, pp. 381-401.
3. *Focus on Success: Using in silico optimisation to achieve an optimal balance of properties.* **Segall, MD, et al.** s.l. : 235, 2006, *Expert Opin. Drug Metab. Toxicol.*, Vol. 2.
4. *Beyond Profiling: Using ADMET models to guide decisions.* **Segall, MD, et al.** 2009, *Chemistry & Biodiversity*, Vol. 6, pp. 2144 - 2151.
5. *Drug Guru: a computer software program for drug design using medicinal chemistry rules.* **Stewart, KD, Shiroda, M and James, CA.** 2006, *Bioorg. Med. Chem.*, Vol. 14, pp. 7011-22.
6. *Evolving molecules using multi-objective optimization: applying to ADME/Tox.* **Ekins, S, Honeycutt, JD and Metz, JT.** 2010, *Drug Discov. Today*, Vol. 15, pp. 451-60.
7. *Rationalizing lead optimization by associating quantitative relevance with molecular structure modification.* **Raymond, JW, Watson, IA and Mahoui, A.** 2009, *J. Chem. Inf. Model.*, Vol. 49, pp. 1952-62.
8. **Burger, A.** *Medicinal Chemistry.* 3rd. San Francisco : John Wiley & Sons Inc, 1970.
9. *Modulation of leukocyte genetic expression by novel purine nucleoside analogues. A new approach to antitumor and antiviral agents.* **Bonnet, PA and Robins, RK.** 1993, *J. Med. Chem.*, Vol. 36, pp. 635-653.
10. *Analogues and derivatives of tenoxicam. 1. Synthesis and antiinflammatory activity of analogues with different residues on the ring nitrogen and the amide nitrogen.* **Binder, D, et al.** 1987, *J. Med. Chem.*, Vol. 30, pp. 678-682.
11. *Discovery of the novel antithrombotic agent 5-chloro-N-((5S)-2-oxo-3-[4-(3-oxomorpholin-4-yl)phenyl]-1,3-oxazolidin-5-yl)methyl)thiophene-2-carboxamide (BAY 59-7939): an oral, direct factor Xa inhibitor.* **Roehrig, S, et al.** 2005, *J. Med. Chem.*, Vol. 48, pp. 5900-5908.
12. *Bioisosterism: A Rational Approach in Drug Design.* **Patani, GA and LaVoie, EJ.** 1996, *Chem. Rev.*, Vol. 96, pp. 3147-3176.
13. *Comparison of some properties of pronethalol and propranolol.* **Black, JW, Duncan, WA and Shanks, RG.** 1965, *Br. J. Pharmacol. Chemother.*, Vol. 25, pp. 577-591.
14. *Synthesis and antiallergy activity of 4-(diarylhydroxymethyl)-1-[3-(aryloxy)propyl]piperidines and structurally related compounds.* **Walsh, DA, Franzysen, SK and Yanni, JM.** 1989, *J. Med. Chem.*, Vol. 32, pp. 105-118.
15. *New dual inhibitors of neutral endopeptidase and angiotensin-converting enzyme: rational design, bioavailability, and pharmacological responses in experimental hypertension.* **Fournié-Zaluski, MC, et al.** 1994, *J. Med. Chem.*, Vol. 37, pp. 1070-1083.
16. *A NEW BIO-ISOSTERE: ALKYL SULPHONAMIDOPHENETHANOLAMINES.* **Larsen, AA and Lish, PM.** 1964, *Nature*, Vol. 203, pp. 1283-1284.

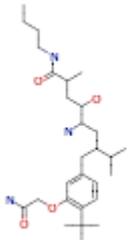
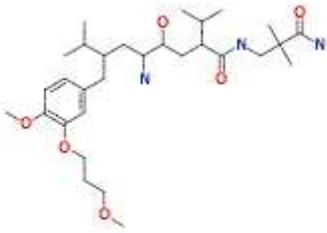
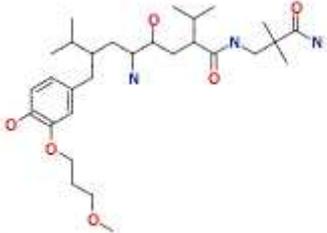
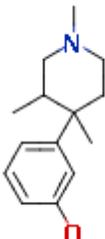
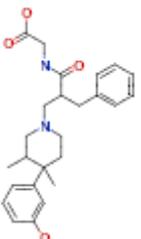
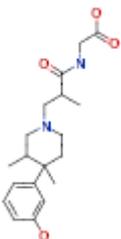
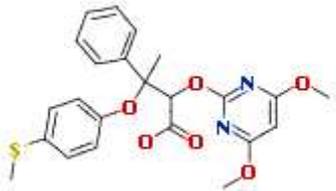
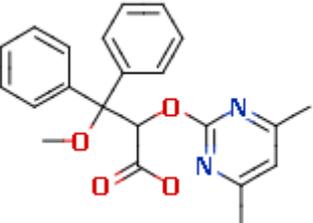
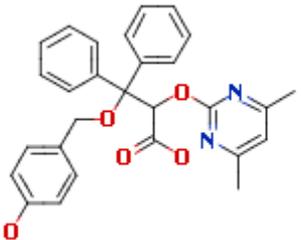
17. *Synthesis and antiviral activity of new anti-HIV amprenavir bioisosteres*. **Rocheblave, L, et al.** 2002, *J. Med. Chem.*, Vol. 45, pp. 3321-3324.
18. *Organic phosphorus compounds. 2. Synthesis and coronary vasodilator activity of (benzothiazolylbenzyl) phosphonate derivatives*. **Yoshino, K, et al.** 1989, *J. Med. Chem.*, Vol. 32, pp. 1528-1532.
19. *Synthesis of antimicrobial agents. 3. Syntheses and antibacterial activities of 7-(4-hydroxypiperazin-1-yl)quinolones*. **Uno, T, et al.** 1990, *J. Med. Chem.*, Vol. 33, pp. 2929-2932.
20. *(S)-3-methyl-5-(1-methyl-2-pyrrolidinyl) isoxazole (ABT 418): a novel cholinergic ligand with cognition-enhancing and anxiolytic activities: I. In vitro characterization*. **Arneric, SP, et al.** 1994, *J. Pharmacol. Exp. Ther.*, Vol. 270, pp. 310-318.
21. *Design, Synthesis, and Anti-inflammatory Properties of Orally Active 4-(Phenylamino)-pyrrolo[2,1-f][1,2,4]triazine p38 α Mitogen-Activated Protein Kinase Inhibitors*. **Hynes Jr., J, et al.** 2008, *J. Med. Chem.*, Vol. 51, pp. 4-16.
22. *Synthesis and evaluation of terbenzimidazoles as topoisomerase I inhibitors*. **Sun, Q, et al.** 1995, *J. Med. Chem.*, Vol. 38, pp. 3638-44.
23. *1,4-Benzodiazepine-2,5-diones as small molecule antagonists of the HDM2-p53 interaction: discovery and SAR*. **Parks, DJ, et al.** 2005, *Bioorg. Med. Chem. Lett.*, Vol. 15, pp. 765-70.
24. *Kinesin spindle protein (KSP) inhibitors. Part 1: The discovery of 3,5-diaryl-4,5-dihydropyrazoles as potent and selective inhibitors of the mitotic kinesin KSP*. **Cox, CD, et al.** 2005, *Bioorg. Med. Chem. Lett.*, Vol. 15, pp. 2041-2045.
25. *SMILES, a chemical language and information system. 1. Introduction to methodology and encoding rules*. **Weininger, D.** 1998, *J. Chem. Inf. Comput. Sci.*, Vol. 28, pp. 31-36.
26. *Computation and Management of Chemical Properties in CACTVS: An extensible Networked Approach toward Modularity and Flexibility*. **Ihlenfeldt, WD, et al.** 1994, *J. Chem. Inf. Comp. Sci.*, Vol. 34, pp. 109-116.
27. **Optibrium**. [Online] [Cited: 3 March 2011.] <http://www.optibrium.com/stardrop>.
28. *Automatic QSAR modeling of ADME properties: blood-brain barrier penetration and aqueous solubility*. **Obrezanova, O, et al.** 2009, *J. Comput. Aided Mol. Des.*, Vol. 22, pp. 431-440.
29. *Gaussian processes: a method for automatic QSAR modeling of ADME properties*. **Obrezanova, O, et al.** 2007, *J. Chem. Inf. Model.*, Vol. 47, pp. 1847-1857.
30. *ChEMBL. An interview with John Overington, team leader, chemogenomics at the European Bioinformatics Institute Outstation of the European Molecular Biology Laboratory (EMBL-EBI). Interview by Wendy A. Warr*. **Overington, J.** 2009, *J. Comput. Aided Mol. Des.*, Vol. 23, pp. 195-198.
31. *DrugBank: a knowledgebase for drugs, drug actions and drug targets*. **Wishart, DS, et al.** 2008, *Nucleic Acids Res.*, Vol. 36(Database issue), pp. D901-6.
32. *An analysis of the binding efficiencies of drugs and their leads in successful drug discovery programs*. **Perola, E.** 2010, *J. Med. Chem.*, Vol. 53, pp. 2986-2997.
33. *RDKit: Cheminformatics and Machine Learning Software*. [Online] [Cited: 2 March 2011.] <http://www.rdkit.org/>.

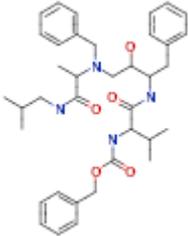
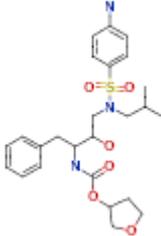
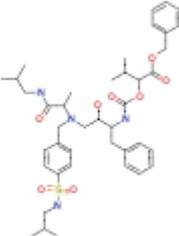
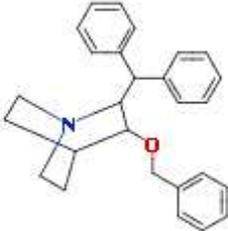
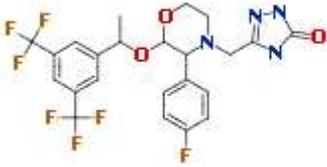
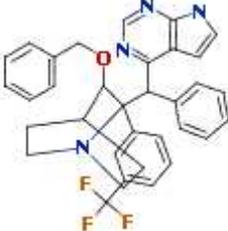
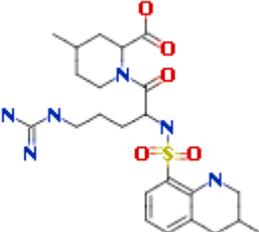
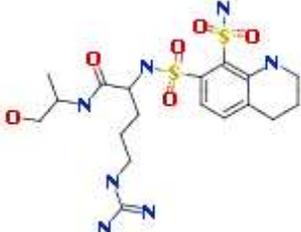
34. **Bolton, E, et al.** PubChem: Integrated Platform of Small Molecules and Biological Activities. *Annual Reports in Computational Chemistry*. Washington DC : American Chemical Society, 2008, Vol. 4, pp. 217-241.

35. *Design and optimisation of selective serotonin re-uptake inhibitors with high synthetic accessibility: part 2.*
Andrews, M, et al. 2009, *Bioorg. Med. Chem. Lett.*, Vol. 19, pp. 5893-5897.

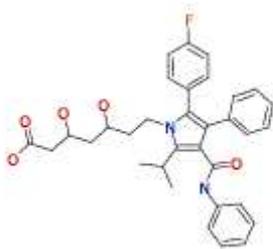
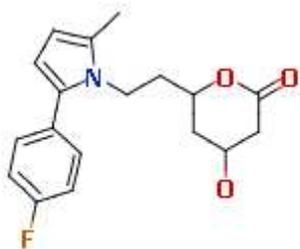
Supporting Information

Table of initial leads, marketed drugs and closest child compounds generated in 5 generations for Perola Set

Drug Name	Initial Lead Structure	Drug Structure	Closest Child Generated	Lead-Drug Similarity	Child-Drug Similarity
ALISKIREN				0.711	0.985
ALVIMOPAN				0.744	0.937
AMBRISENTAN				0.659	0.931

AMPRENAVIR				0.621	0.737
APREPITANT				0.530	0.720
ARGATROBAN				0.540	0.862
ATAZANAVIR				0.755	0.842

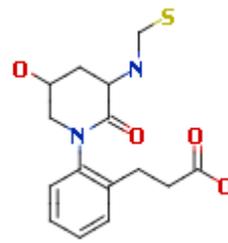
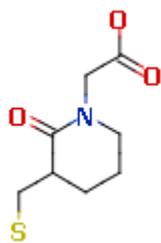
ATORVASTATIN



0.644

0.925

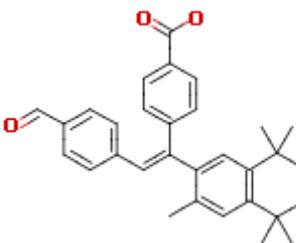
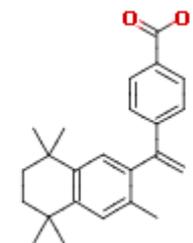
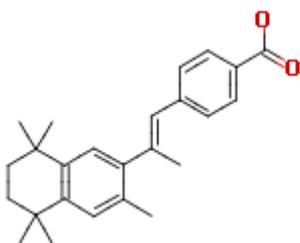
BENAZEPRILAT



0.468

0.786

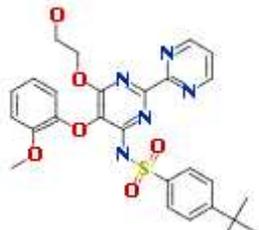
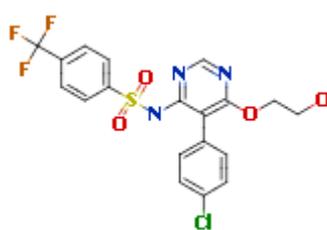
BEXAROTENE



0.743

0.911

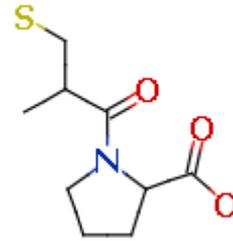
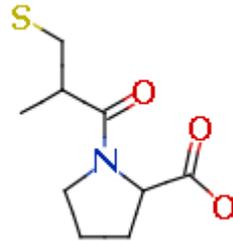
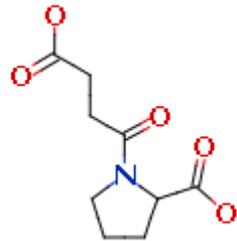
BOSENTAN



0.724

0.825

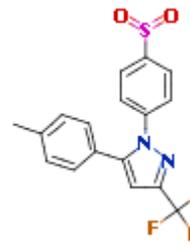
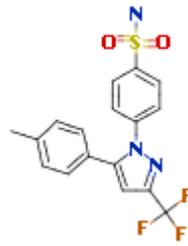
CAPTOPRIL



0.791

1.000

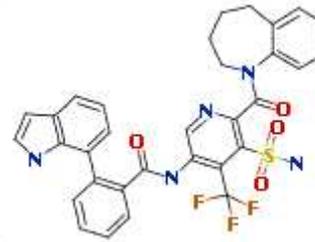
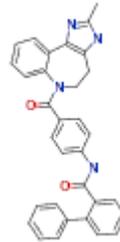
CELECOXIB



0.905

0.967

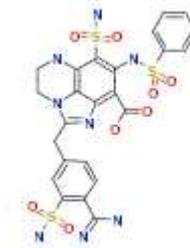
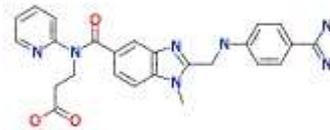
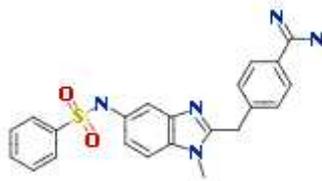
CONIVAPTAN



0.643

0.851

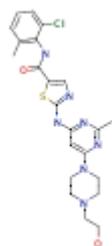
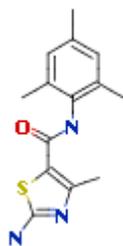
DABIGATRAN



0.639

0.794

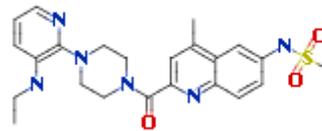
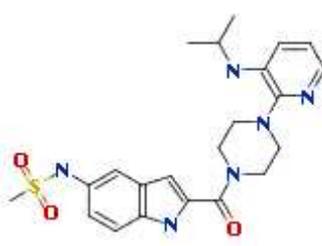
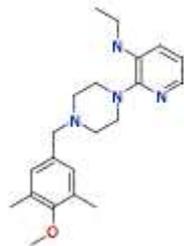
DASATINIB



0.624

0.793

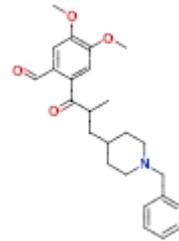
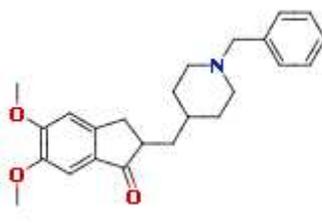
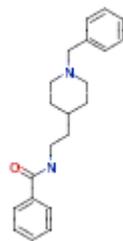
DELAVIRDINE



0.650

0.911

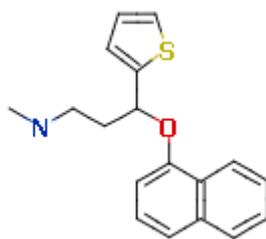
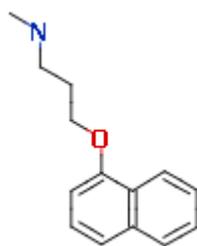
DONEPEZIL



0.493

0.920

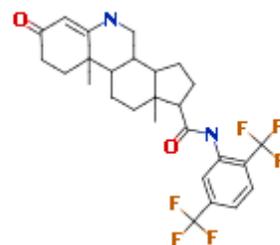
DULOXETINE



0.497

1.000

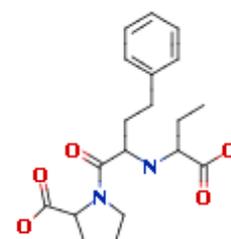
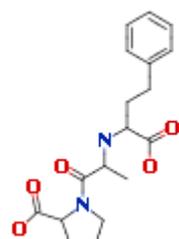
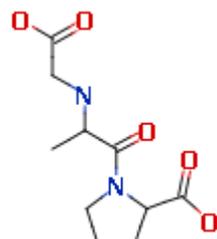
DUTASTERIDE



0.549

0.838

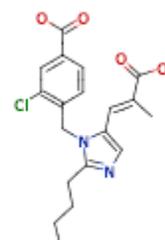
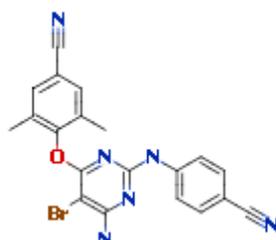
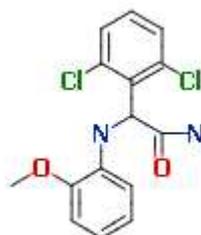
ENALAPRILAT



0.787

0.921

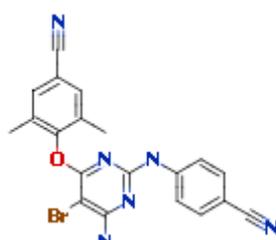
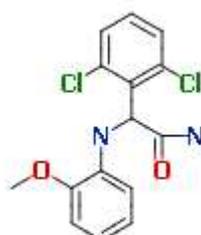
EPROSARTAN



0.659

0.821

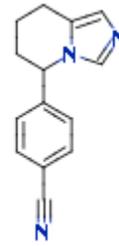
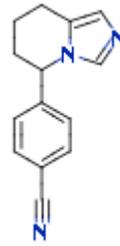
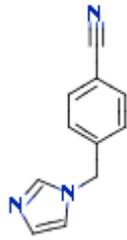
ETRAVIRINE



0.473

0.660

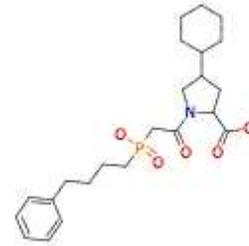
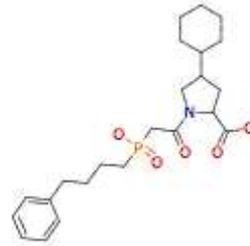
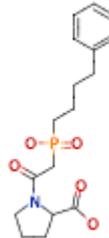
FADROZOLE



0.488

1.000

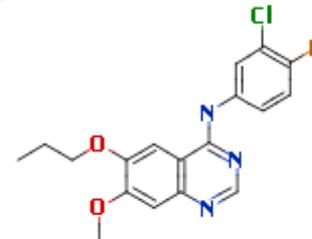
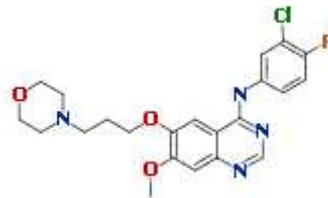
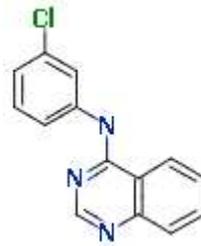
FOSINOPRILAT



0.878

1.000

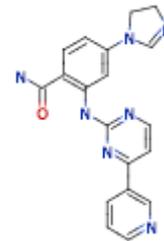
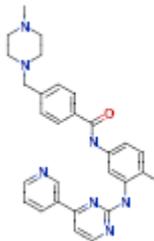
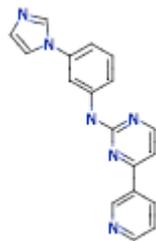
GEFITINIB



0.651

0.952

IMATINIB



0.650

0.786

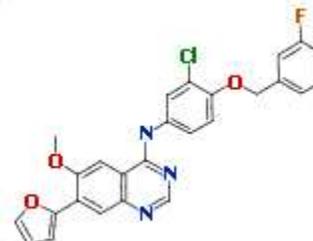
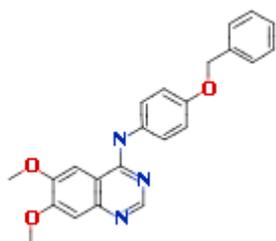
INDINAVIR



0.675

0.839

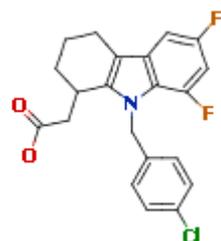
LAPATINIB



0.674

0.833

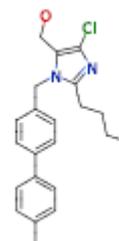
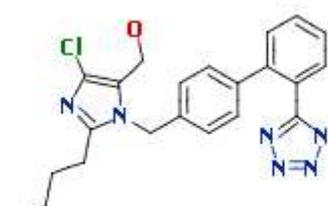
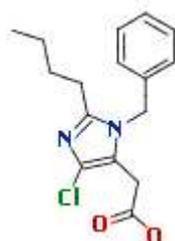
LAROPIPRANT



0.893

0.938

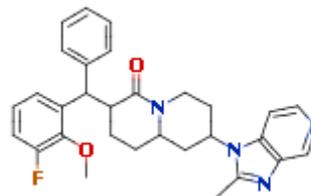
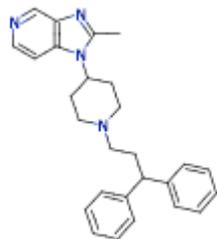
LOSARTAN



0.733

0.866

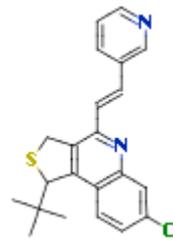
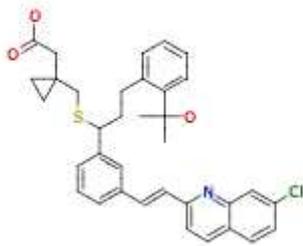
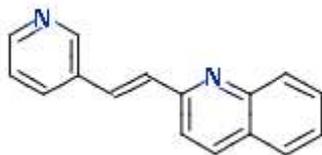
MARAVIROC



0.604

0.729

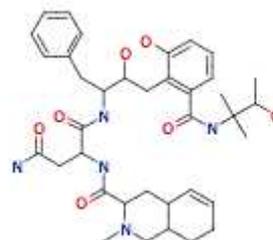
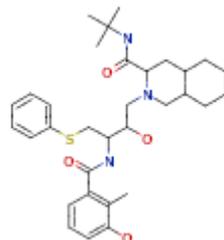
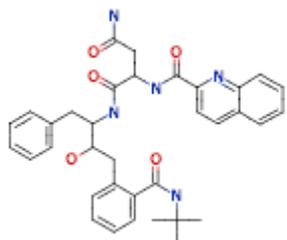
MONTELUKAST



0.461

0.732

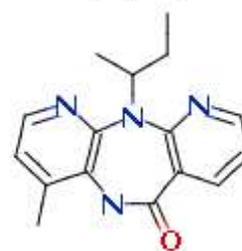
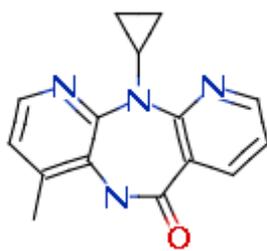
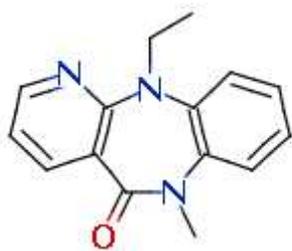
NELFINAVIR



0.676

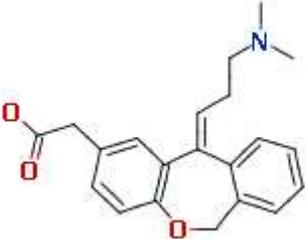
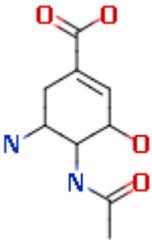
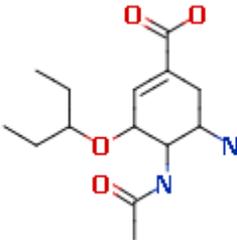
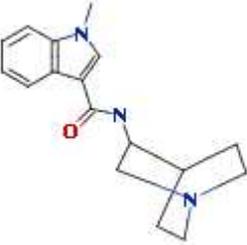
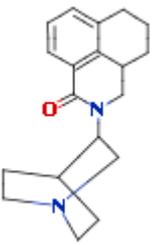
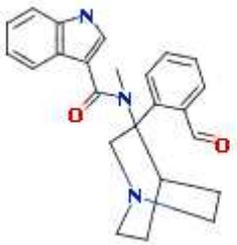
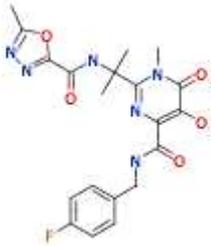
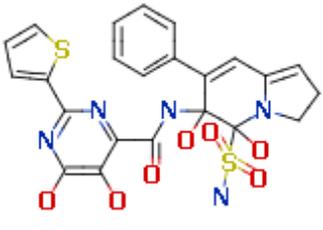
0.768

NEVIRAPINE

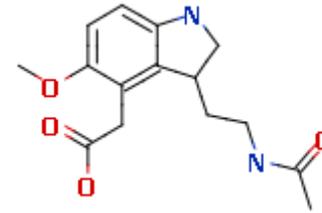
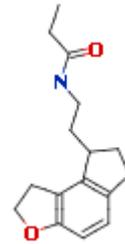


0.823

0.990

OLOPATADINE				0.634	0.846
OSELTAMIVIR CARBOXYLATE				0.835	1.000
PALONOSETRON				0.631	0.771
RALTEGRAVIR				0.682	0.859

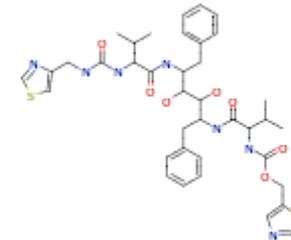
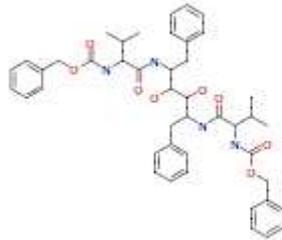
RAMELTEON



0.416

0.645

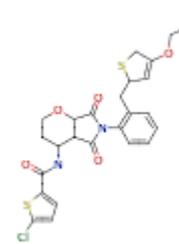
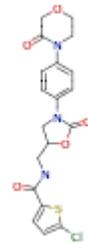
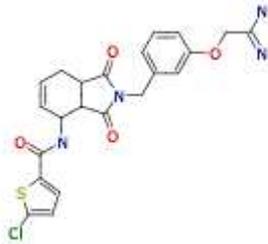
RITONAVIR



0.714

0.884

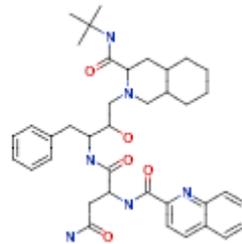
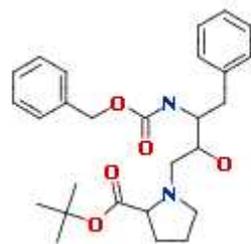
RIVAROXABAN



0.619

0.688

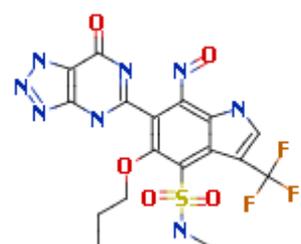
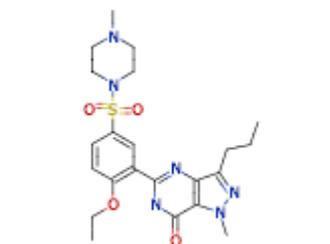
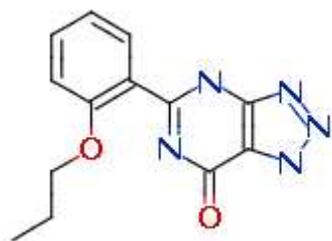
SAQUINAVIR



0.614

0.807

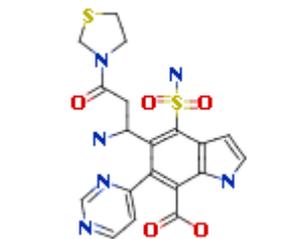
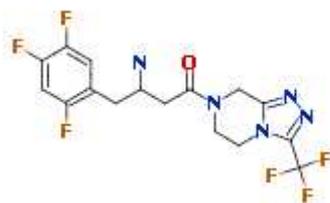
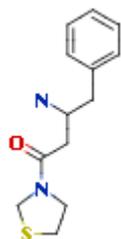
SILDENAFIL



0.653

0.882

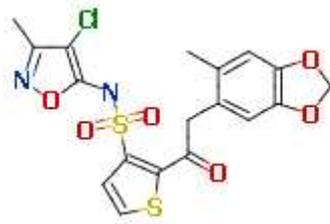
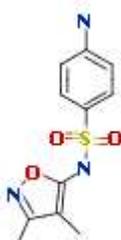
SITAGLIPTIN



0.402

0.709

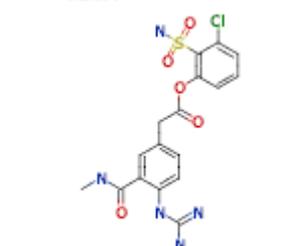
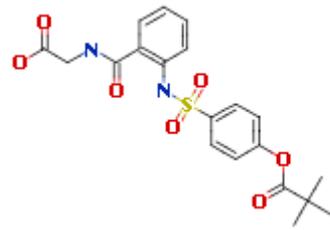
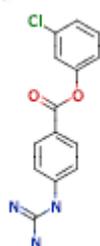
SITAXENTAN



0.570

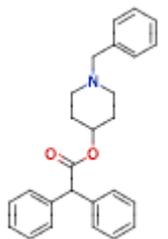
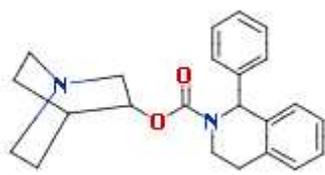
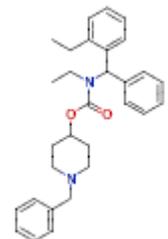
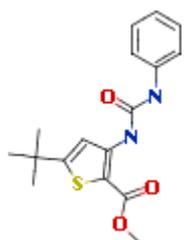
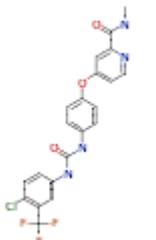
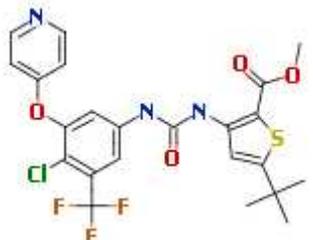
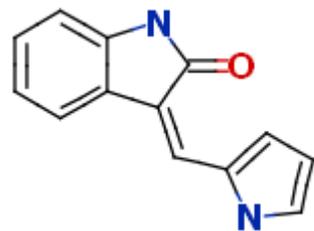
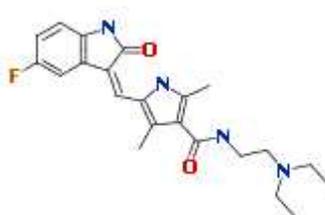
0.842

SIVELESTAT

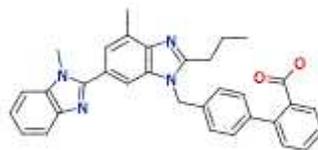
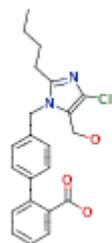


0.380

0.651

SOLIFENACIN				0.446	0.795
SORAFENIB				0.479	0.634
SUNITINIB				0.721	0.957
TADALAFIL				0.502	0.919

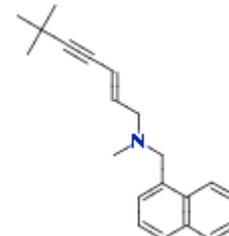
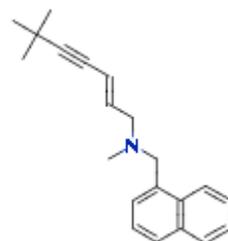
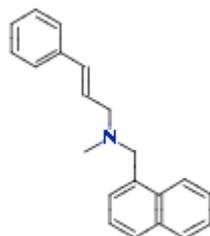
TELMISARTAN



0.718

0.845

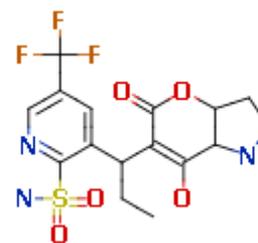
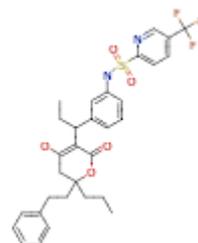
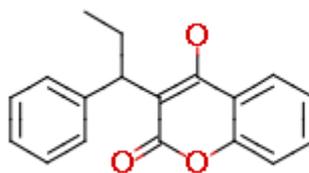
TERBINAFINE



0.757

1.000

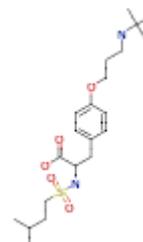
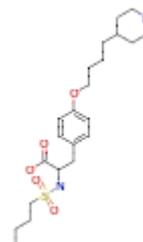
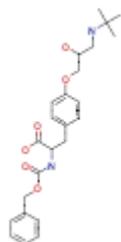
TIPRANAVIR



0.601

0.879

TIROFIBAN



0.528

0.883

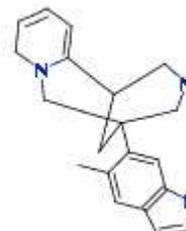
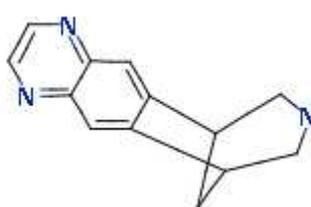
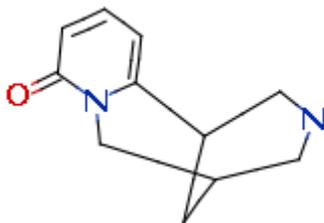
TOPOTECAN



0.985

1.000

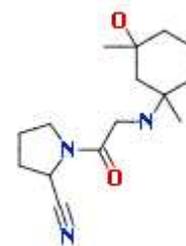
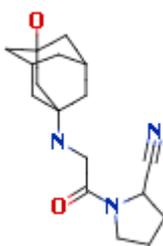
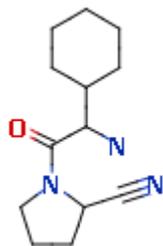
VARENICLINE



0.455

0.604

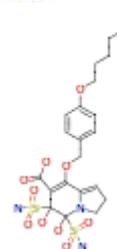
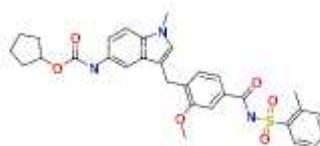
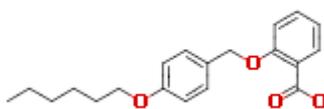
VILDAGLIPTIN



0.592

0.970

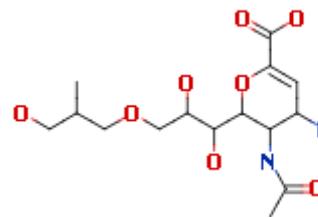
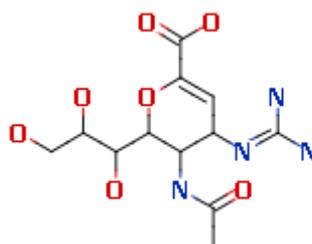
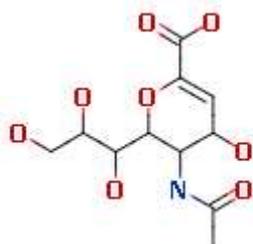
ZAFIRLUKAST



0.373

0.838

ZANAMIVIR



0.797

0.908

Average

0.636

0.853