

AutoDesigner, a *De Novo* Design Algorithm for Rapidly Exploring Large Chemical Space for Lead Optimization: Application to the Design and Synthesis of D-Amino Acid Oxidase Inhibitors

Pieter H. Bos,[‡] Evelyne M. Houang,[‡] Fabio Ranalli,[‡] Abba E. Leffler,[‡] Nicholas A. Boyles, Volker A. Eyrich, Yuval Luria, Dana Katz, Haifeng Tang, Robert Abel, and Sathesh Bhat*

[‡]Authors contributed equally to this work

*Corresponding author; Email: sathesh.bhat@schrodinger.com

Schrödinger, Inc., 1540 Broadway, 24th floor, New York, NY 10036, United States

Abstract

The lead optimization stage of a drug discovery program generally involves the design, synthesis and assaying of hundreds to thousands of compounds. The design phase is usually carried out via traditional medicinal chemistry approaches and/or structure based drug design (SBDD) when suitable structural information is available. Two of the major limitations of this approach are (1) difficulty in rapidly designing potent molecules that adhere to myriad project criteria, or the multiparameter optimization (MPO) problem, and (2) the relatively small number of molecules explored compared to the vast size of chemical space. To address these limitations we have developed AutoDesigner, a *de novo* design algorithm. AutoDesigner employs a cloud-native, multi-stage search algorithm to carry out successive rounds of chemical space exploration and filtering. Millions to billions of virtual molecules are explored and optimized while adhering to a customizable set of project criteria such as physicochemical properties and potency. Additionally, the algorithm only requires a single ligand with measurable affinity and a putative binding model as a starting point, making it amenable to the early stages

of a SBDD project where limited data is available. To assess the effectiveness of AutoDesigner, we applied it to the design of novel inhibitors of D-amino acid oxidase (DAO), a target for the treatment of schizophrenia. AutoDesigner was able to generate and efficiently explore over 1 billion molecules to successfully address a variety of project goals. The compounds generated by AutoDesigner that were synthesized and assayed (1) simultaneously met not only physicochemical criteria, clearance and central nervous system (CNS) penetration (K_{p,uu}) cutoffs, but also potency thresholds; (2) fully utilize structural data to discover and explore novel interactions and a previously unexplored subpocket in the DAO active site. The reported data demonstrate that AutoDesigner can play a key role in accelerating the discovery of novel, potent chemical matter within the constraints of a given drug discovery lead optimization campaign.

Introduction

The aim of most current small-molecule drug discovery efforts is to develop several diverse, patentable series of compounds that potently bind to a target protein while also having favorable properties such as solubility, permeability, binding specificity and metabolic stability as specified in a Target Product Profile (TPP).¹ It is well documented that discovery teams find this multiparameter optimization problem challenging, often delivering high molecular weight (MW), lipophilic compounds that consequently possess poor pharmacokinetic (PK) profiles, a phenomenon described as “molecular obesity”.² As a result, there has been considerable interest recently in using Artificial Intelligence (AI) to generate compounds *de novo* under the premise that algorithms will not be subject to such human behavioral biases and thus deliver higher quality compounds.³ While a number of these algorithms have reported success on retrospective^{4,5} and prospective tests,⁶ there are comparatively fewer reports of their performance in the setting of an active drug discovery program.

We hypothesized that a *de novo* compound design algorithm that emulates the design approaches of medicinal chemists, searches chemical space on a vast scale, and remains constrained to a chemical property space desirable for good PK would be able to produce diverse, drug-like, and potent compounds. To test this hypothesis, we have created a *de novo* design algorithm, AutoDesigner, that is based on a number of ligand design strategies that medicinal chemists commonly use during the Hit-to-Lead (H2L) and Lead Optimization (LO) discovery stages, does not require a large amount of prior experimental data, and is inexpensive to run relative to the cost of synthesizing compounds. The algorithm uses a combination of matched molecular pair (MMP) transforms, reaction-based enumeration, and R-group decoration to explore chemical space on an ultra-large scale, routinely on the order of 100 million to >1 billion molecules, usually with runtimes on the order of 24-48 hours. These iterative generative cycles of exploration are coupled with an extensive filtering cascade that evaluates the molecules against a large number of endpoints such as physicochemical properties, complexity, and the ability to dock in the receptor. The end result is a large pool of diverse, drug-like compounds that can fit into the target protein's binding pocket, thoroughly explore available vectors in a synthetically tractable fashion, and whose potency can readily be evaluated using structure-based computational methods like Free-Energy Perturbation (FEP+).⁷

As a challenging real-world, prospective test case for the AutoDesigner algorithm, we chose the development of inhibitors of D-amino acid oxidase (DAO). Studies have shown that inhibiting DAO to increase levels of D-serine, an endogenous agonist of the glycine modulatory site of the N-methyl D-aspartate receptor (NMDAR), could be a valid therapeutic approach for treating schizophrenia.^{8,9,10,11} A variety of DAO inhibitors binding the catalytic site have been identified^{12,13} and several public and in-house high-quality co-crystal structures have been made available to use as the starting point for an AutoDesigner effort (Figure 1). Full details of our DAO discovery program can be found elsewhere.¹⁴ In this work, we demonstrate how

AutoDesigner was used to develop novel DAO inhibitors that met a challenging TPP and suggest more broadly how it might find utility in a variety of discovery programs.




Figure 1. Crystal structure showing the binding mode and interactions of compound 1 (green) in the DAO binding site with coenzyme flavin adenine dinucleotide (FAD) (magenta).

Methods

The AutoDesigner algorithm consists of a multi-stage search algorithm that utilizes cloud computing to carry out successive rounds of chemical space exploration and filtering. This *de novo* design algorithm only requires a single ligand with measurable binding affinity and a

putative binding model, making it amenable to the early stages of a SBDD project where limited data is available. Over the course of several generative stages and subsequent filtering stages, millions to billions of molecules are explored and evaluated against a customizable set of project criteria. A general overview of the AutoDesigner algorithm is depicted in Figure 2.

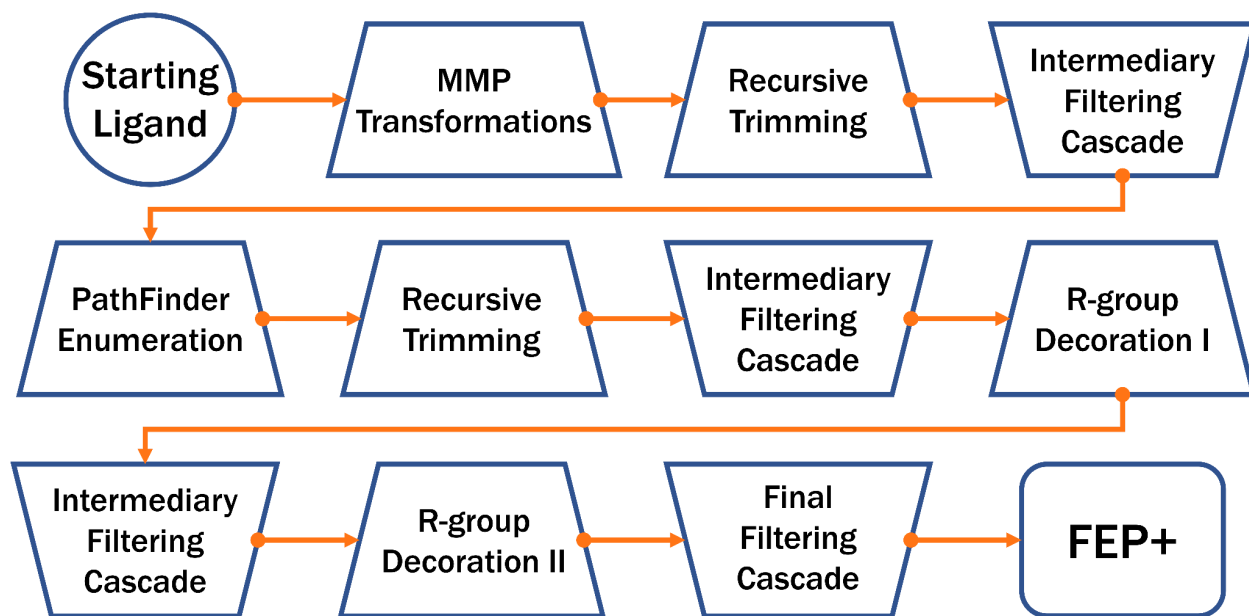


Figure 2. General overview of the AutoDesigner algorithm. Generative stages are trapezoid-shaped, and filtering stages are depicted as inverted trapezoids.

The algorithm utilizes four different methods to explore chemical space: 1) Matched molecular pair (MMP) transformations; 2) PathFinder reaction-based enumeration; 3) Recursive Trimming; 4) R-group decoration. These stages were designed to mimic common ideation strategies employed by medicinal chemists. The outputs from the generative stages are funneled into the intermediary filtering cascades, and after filtering the remaining molecules are used as inputs for the next generative stage. Ultimately, all compounds that pass the intermediary filtering cascades are combined, deduplicated, and passed through the final filtering cascade to return the final AutoDesigner output compounds. The docked poses

generated in the AutoDesigner algorithm are then used as starting points for FEP+ potency predictions. If an AutoDesigner run generates tens of thousands of compounds, then the processing of all compounds can be sped up effectively using an active learning approach based on FEP+ predictions.^{15,16} In Active-Learning FEP (AL-FEP), single-edge FEP+ predictions (utilizing one central reference compound) from a random subset are used to train a machine learning (ML) model, and the resulting ML model is used to score the remaining compounds in the library. Based on the ML scores, the next set of compounds is selected for single-edge FEP+ profiling, and a new ML model is trained on the cumulative predictions. This iterative cycle is repeated until all promising compounds have been retrieved from the library.¹⁷

Matched Molecular Pair Transformations

Matched molecular pairs (MMP) were generated by fragmentation of the molecules from the PubChem¹⁸ and ChEMBL¹⁹ databases using a maximum heavy atom difference of 8 atoms and a maximum of 2 cuts.²⁰ The resulting outputs were converted to their respective reaction SMARTS to be applied as transformations. In total, 291 million and 26 million unique transformations were generated from the PubChem and ChEMBL databases, respectively. A more detailed analysis of the extracted transforms can be found in the Supplementary Material (S1).

Before applying the set of transforms, a set of input compounds is generated from a single starting ligand to introduce additional chemical matter to maximize the number of applicable transformations. The full set of 300 million transforms is applied to these input molecules and only ligands that maintain the predefined immutable core are kept. The resulting output is trimmed recursively, as described below, to generate additional ligands. The ligands are filtered using an intermediary filtering cascade, as described below, and used as inputs in the following generative stage.

Reaction-based Enumeration

All the passing compounds from the MMP transformation stage are used as inputs for the reaction-based enumeration stage. This generative stage performs a retrosynthetic analysis of the input molecules, up to a configurable depth, followed by combinatorial synthesis of the resulting reaction pathways using our PathFinder tool.¹⁵ All routes for each input ligand are enumerated using commercially available reactant libraries. The reactant containing the immutable region, or common core, of the ligand is kept constant while the other reactant(s) are varied. The diverse nature of the input ligands, originating from the matched molecular pair transformation stage, generates a large variety of possible routes, and ensures a diverse exploration of synthetically-accessible chemical space. After deduplication, the ligands are recursively trimmed, as described below, and the resulting output is filtered in the intermediary filtering cascade.

Recursive trimming

Before the outputs resulting from both the MMP transformation stage as well as the reaction-based enumeration stage are progressed into the filtering cascade, a recursive trimming algorithm is applied to generate additional ligands. This algorithm breaks predefined single bonds in a recursive fashion, while at the same time keeping the bonds of the immutable core intact, leading to a large number of additional ligands generated from all possible relevant permutations (Figure 3). Rules are applied to prevent trimming of bonds that would lead to undesirable outputs (e.g., aromatic bonds).




Figure 3. Example of the recursive trimming algorithm applied to the boxed compound, resulting in seven additional design ideas. The immutable region is highlighted in blue.

R-group Decoration

In the final two generative stages, all ligands that survived the intermediary filtering cascades following the MMP transformation stage and the reaction-based enumeration stage are combined. The resulting ligands are then decorated using a curated library of R-groups. Each input is decorated at every available site excluding the fixed core. After the first decoration round, the resulting products are funneled through the intermediary filtering cascade, and the remaining compounds are subjected to another round of R-group decoration. Only ligands that can accommodate the extra atoms from the R-group are decorated. Two rounds of decorations ensure a broad exploration of the available chemical space. After the second round of decoration, the resulting products are combined with all of the ligands that passed the intermediary filtering cascades of the first three generative steps (MMP transformation,

reaction-based enumeration, and R-group decoration). After deduplication of this set, the ligands are funneled into the final filtering cascade.

Filtering Stages

Of equal importance to generating ligands and exploring chemical space is the ability to effectively filter these large numbers of compounds so that only the compounds that match the desired profile are retained. The AutoDesigner algorithm contains a progressive filtering cascade that employs a range of different filtering techniques. The order of the filtering steps is based on their computational efficiency and there are two different types of filtering cascades. The intermediary filtering stage uses fewer filters and is set up to let compounds through that could potentially be rescued in later stages, e.g., compounds are not filtered based on the number of hydrogen bond donors as these can be added or removed in the decoration stage. The goal of the intermediary filtering steps is to remove undesired compounds efficiently without removing promising compounds. The final filtering stage uses all available filters and the strictest settings so that only compounds that fit all the desired criteria are retained. Due to the ultra-large-scale exploration of chemical space, strict filtering is both advantageous and necessary in order to generate the best possible compounds.

In the first step of the filtering cascade, compounds are filtered based on >20 physicochemical properties (e.g., MW, TPSA, logP etc.) and their overall complexity (e.g., number of chiral centers, spiro atoms etc.; see Supplementary Material S2-S4 for additional details). Additionally, compounds are also analyzed for their overall complexity compared to the starting ligand by looking at structural features, e.g., the number of chiral centers, spiro atoms, and all-carbon quaternary centers. The complexity filter reduces the number of compounds that would be too challenging to synthesize, independent from the complexity of the core and the overall physicochemical property space.

The compounds that pass all of these criteria are progressed to the next filtering stage in which the ligands are filtered using a proprietary set of ~1,500 SMARTS patterns to remove known and potential chemical liabilities, PAINS offenders, and other undesired chemotypes. Every AutoDesigner run explores different targets and areas of chemical space, and therefore the outputs of the runs are periodically reviewed by our in-house team of experienced medicinal chemists to flag compounds that are either chemically unstable, contain medicinal chemistry liabilities, or for which the synthesis is unprecedented. This resulted in continuous improvement of the algorithm and the SMARTS patterns used for substructure filtering over an eight month period. Although originally as many as 20% of compounds were manually flagged, in later runs less than 1% of the AutoDesigner-generated compounds were flagged by visual inspection (Figure 4).

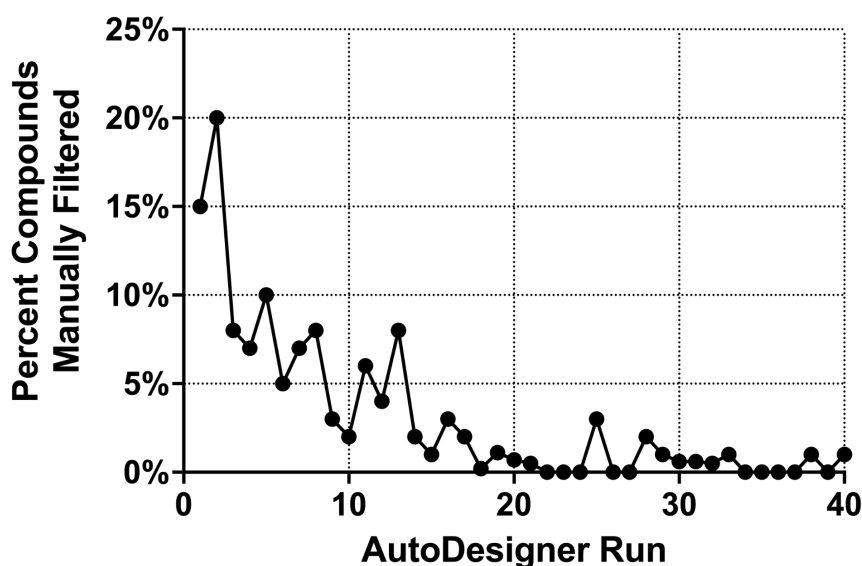


Figure 4. Percentage of AutoDesigner compounds manually filtered out by medicinal chemists over time.

In the next step of the filtering process the ligands are evaluated for their FEP+ amenability, which is a key property required to ensure accurate FEP+ predictions.²¹ Each ligand that makes it to this part of the filtering cascade is compared with relevant FEP+ references that are provided as part of the setup, and the heavy atom perturbation between ligand and reference is calculated. Based on FEP+ performance for the area of interest, a maximum heavy atom perturbation is set and the ligands are filtered accordingly. The ligands are then prepared for docking using the LigPrep module from Maestro (Schrödinger, Inc.: New York, NY, USA), and only ligands with a total and absolute charge equal to the desired charge state used in the docking and FEP+ models are retained. The state penalties (i.e., relative energies) of the different ionization and tautomeric states for the individual ligands are calculated and ligands with a high state penalty are removed to reduce false positives and irrelevant ligand states.

The resulting library of ligands is then subjected to a number of optional filters that are fully configurable. If the ligands contain a titratable center in the newly designed R-group, then a pKa filter can be used to filter for ligands in the desired pKa range. This allows for precise tuning of the basicity and associated ADMET properties. Additionally, compounds can be filtered on their calculated logD values.

Next, the filtered and prepared 3D structures are funneled into the docking stage. In this stage the ligands are docked in the receptor using a maximum-common-substructure-based docking model combined with optional hydrogen bond constraints.²² Ligands that do not fit in the binding site and/or do not match the hydrogen bond constraints are discarded and the ligands that docked successfully proceed to the next stage. Following the docking stage there is an optional filtering step that analyzes the docked poses based on a pharmacophore model using a Phase hypothesis.^{23,24} All the final ligands are then combined, deduplicated, and written to a ligand library file containing the docked poses.

The output from the AutoDesigner algorithm can be triaged further by applying additional filters, e.g., machine learning models trained on relevant endpoints, and the final set of docked poses are used directly as input poses for active-learning FEP.^{15,16}

Relative Binding Free Energy Calculation protocol

Protein-ligand bound X-ray structures were obtained via collaboration with Takeda Pharmaceuticals and processed with the Protein Preparation Wizard in Maestro²⁵ using default settings. All ligand structures were processed using LigPrep to enumerate all stereoisomers, and protonation states were assigned using Epik.²⁶ Possible tautomer and charge states were evaluated using Jaguar pKa²⁷ and corrected for in the Free Energy calculations.²⁸ OPLS3e²⁹ force field torsion parameters for each ligand were generated using the Force Field Builder. Prospective free energy calculations were run using Schrödinger FEP+⁷ with a simulation time of 20 ns, 24 λ -windows, and grand canonical Monte Carlo (GCMC) enhanced water sampling.³⁰

Generation of AutoQSAR machine learning models for K_{p,uu} and Clearance

Machine learning models trained on experimental *in vivo* mouse PO K_{p,uu} (brain/plasma) and mouse IV plasma clearance (Cl_p) project endpoints were generated using AutoQSAR, an automated QSAR model builder application.³¹ In brief, all experimental data measured at that particular point in time during the course of the project (45 K_{p,uu} and 54 Cl_p data points at initial model build), were separated into a training set and a holdout set using an 80/20 time-split. The top scoring models were then used prospectively as part of the AutoDesigner algorithm. Each model was retrained over time as more experimental data was collected.

Implementation on a cloud computing platform

The AutoDesigner algorithm has been implemented as a multi-stage algorithm which is, by design, very bursty in terms of compute resource utilization: it can typically require a compute

capacity in the 2K-150K CPU cores range, with peaks that can last for relatively short amounts of time, at different stages of the algorithm. Therefore, we designed and implemented a bespoke infrastructure, deployed on a public cloud platform, that was optimized to run the AutoDesigner algorithm in a highly efficient manner. Such infrastructure has the ability to automatically expand and shrink its compute capacity according to the requirements of each and every stage. Having a fully elastic cluster allows for optimization of both cost and performance. On average an end-to-end AutoDesigner run takes 24 - 48 hours. To prevent unnecessary workloads, the AutoDesigner algorithm always executes inline compound deduplication at the end of each and every generative stage. To remove all possible duplicates, the AutoDesigner algorithm executes an across-the-board deduplication stage after each series of generative stages. The deduplication stage merges all generated compounds into a single dataset, and removes duplicates from the entire dataset. We take advantage of a cloud-based, fully serverless data warehouse to perform this task. By ingesting the dataset into a table, and implementing the deduplication process as SQL-like code, we deduplicate from tens to hundreds of billions of compounds from a few minutes to a few hours in rare cases.


Finally, we built automation code that can automatically provision the exact amount of required compute resources. The high performance computing (HPC) part relies on various compute node sizes to run computations. The memory to CPU ratio is 2 GB of RAM per CPU core. Based on the job requirements, one or more compute nodes will be provisioned on-demand, each one with the proper specifications, ranging from a lower bound of 8 CPUs / 16 GB per node, up to 64 CPUs / 128 GB. The cloud infrastructure operates with a 30 CPUs / 60 GB compute node size. The data warehouse part, instead, is completely serverless and autoscaled by design, and therefore does not require any type of pre-configuration other than creating the schema.

Results and Discussion

To efficiently explore relevant chemical space, it is critical to set the physicochemical property space so that it aligns with the overall project goals. In this case, the ranges of the physicochemical properties were set to match the desired CNS property space. A fine balance for these property ranges is required: too narrow, and the diversity and number of compounds generated is limited; too wide, and the exploration is costly and inefficient. The optimal physicochemical property space, along with a robust filtering cascade, allows for efficient and thorough exploration of millions of compounds while at the same time minimizing computational time spent on exploring irrelevant chemical matter. Here, we describe how AutoDesigner solved five program challenges by performing ultra-large-scale design within the property space defined by the project team (Table S2b).

Generation of DAO inhibitors with AutoDesigner

Statistics for three AutoDesigner campaigns carried out during this program are shown in Table 2, on a per-stage basis. For the AutoDesigner run based on compound **1**, a total of 761M ligands were explored, and ultimately 515K compounds survived to the end of the process. For the run based on compound **5**, a total of 353M ligands were explored and 15K passed through all filters. Finally, for the AutoDesigner campaign launched on the basis of compound **7**, although 199M ligands were explored, only 156 were retained.

Table 2: Number of compounds generated during different AutoDesigner stages.^{a,b,c}

Input Ligand	MMP	Trim 1	PF	Trim 2	Deco 1	Deco 2	Total ^d	Final Output
1	157M	24M	466M	24M	44M	45M	761M	515K
5	101M	121M	234M	2M	2M	2M	353M	15K
7	161M	13M	23M	458K	348K	62K	199M	156

^a Immutable region/common core of the starting ligands highlighted in blue. ^b Abbreviations used for generative stages: MMP: matched molecular pair transformations; Trim: recursive trimming; PF: PathFinder reaction-based enumeration; Deco: R-group decoration. ^c Additional details for the individual runs can be found in the Supplementary Material (S2-4). ^d Total number of unique compounds explored

In addition, for each of the three runs the CNS MPO³² was computed for all surviving AutoDesigner compounds (Figure 5). For all three runs, the distributions were fairly tight and centered around median CNS MPO values of roughly 5. This value is greater than the cutoff of 4 which is generally considered a “good” score.³² While some compounds did have CNS MPO values below this value, they represented a minority overall. Taken together, these data suggest that AutoDesigner produced compounds of high quality as measured by the CNS MPO metric, which was calculated subsequent to the AutoDesigner runs and not during the design process.

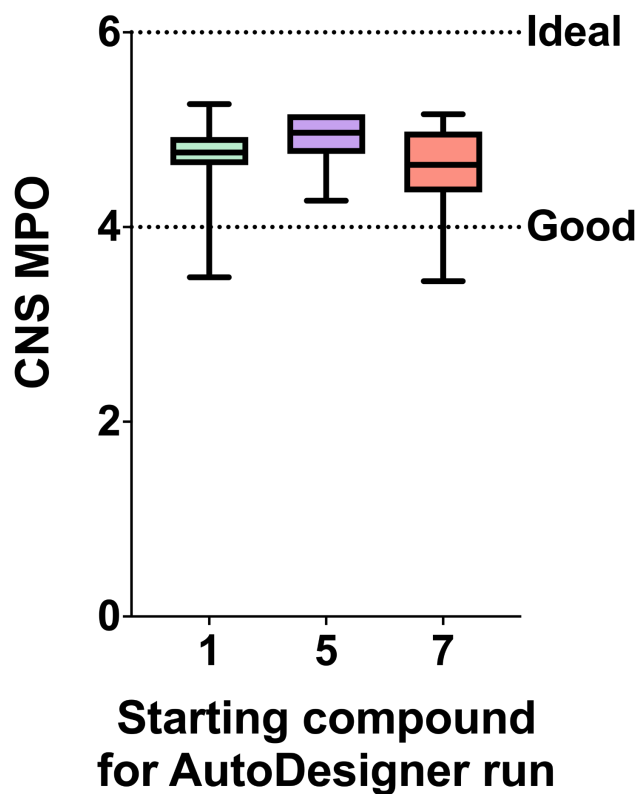



Figure 5: Boxplots of the CNS MPO scores of compounds produced from AutoDesigner campaigns starting from Compounds 1, 5, and 7.

Finally, Scheme 1 exemplifies compounds used as starting points for the different AutoDesigner runs (compounds 1, 5, and 7), and selected compounds generated by AutoDesigner that were predicted to be potent by FEP+ prospectively and subsequently synthesized (compounds 2, 3, 4, 6, and 8).



Scheme 1. Structures of the DAO inhibitors and associated hDAO pIC_{50} ; FEP+ predicted pIC_{50} in parentheses. ^aMeasured pIC_{50} of the racemate.

Application of AutoDesigner to Identify Compounds with Moderate Lipophilicity

One of the key parameters of an efficacious DAO inhibitor is adequate CNS exposure, which in part requires molecules with a moderate lipophilicity.³³ For this reason we decided to use AutoDesigner to explore the relevant physicochemical property space and generate compounds that introduce some additional polarity to offset the higher hydrophobicity of the tail region, while still maintaining other favorable properties of the starting ligand. Compound **1**, consisting of a pyrazine dione head piece, a 2-carbon linker, and a trifluoromethyl phenyl tail piece, was used as the starting point. The AutoDesigner algorithm was used to explore the SAR of the aromatic tail region, and the 2-carbon linker and pyrazine dione were kept constant. The goal was to explore whether additional polarity could be introduced in the tail region without impacting other favorable properties and without a dramatic reduction in enzymatic potency.

The algorithm consists of several expansion stages that generate hundreds of millions of compounds and thoroughly explore the desired property space. Each expansion stage is followed by a filtering cascade that filters ligands based on the desired physicochemical property space, removes undesired and/or unstable chemotypes, and ensures FEP+ amenability. Additionally, only the ligands that docked successfully were promoted to the next stage (see Figure 2 for the general algorithm).




Figure 6. Progression of a compound through AutoDesigner (A) Starting compound **1** (B) MMP Transformation stage (C) PathFinder reaction-based enumeration stage (D) Decoration stage; compound **2**.

The filtering stage is crucial to the algorithm, as it ensures that only compounds that adhere to the predefined project criteria, do not introduce chemical liabilities, and fit within the binding site, propagate to later stages. Following the PathFinder reaction-based enumeration stage, and the subsequent filtering cascade, the ligands are passed on to the decoration stage, which decorates the designed ligands with R-groups (up to 15 heavy atoms). This decoration

stage increases compound diversity and further explores chemical space. It also attempts to capture the “magic-methyl” effect.³⁴ In the example highlighted in Figure 6, the matched molecular pair (MMP) transformation stage converts the phenyl-CF₃ to an indole (Figure 6B), which passes all filters and fits in the binding site. In the next generative stage, the PathFinder reaction-based enumeration stage, the indole is converted into an indazole (Figure 6C). The decoration stage adds a CF₃ group to the indazole-containing ligand (Figure 6C) generated in the PathFinder stage, and compound **2**, together with hundreds of millions of other compounds, enter the final, most comprehensive, filtering cascade. In the end, compound **2** (Figure 6D) and 515,000 other diverse compounds passed this final filtering cascade, confirming their FEP+ amenability, desired physicochemical property space, chemical stability, and ability to fit into the protein binding pocket.

The most promising compounds that passed the final filtering cascade were evaluated on their CNS MPO score, FEP+ predicted potency, and predicted LLE. Based on these parameters compound **2** was one of the analogs that was synthesized, assayed, and confirmed as a potent inhibitor of DAO with a pIC₅₀ = 6.19, and an improved CNS MPO score (Table 3). In line with the specified design criteria, the final compound has a higher TPSA than the starting compound, while retaining potency.

Table 3. Properties of starting and AutoDesigner compounds.

Property	1	2
hDAO pIC ₅₀	6.51	6.19
Pfizer CNS MPO ³²	4.57	4.84
MW	284.24 Da	324.26 Da
Fraction sp ³	0.23	0.21
TPSA	66 Å ²	94 Å ²
AlogP	2.30	2.04
HBA	2	3
HBD	2	3

Rapid Design of Compounds to Accommodate Differences in Protein Conformation

SBDD relies heavily on high resolution crystal structures of the target protein in order to design potent ligands. Initially, an unpublished proprietary X-ray structure was used to profile potential ligands with FEP+. Later on in the project we obtained a crystal structure of the protein bound to compound **1** and two significant differences were observed: (1) His217 moves slightly into the pocket, and (2) Ile223 swings out thereby creating a larger subpocket able to accommodate larger ligands. AutoDesigner was employed to quickly explore the binding site of the new crystal structure.

Simply changing the docking model in the algorithm allows AutoDesigner to rapidly probe the larger binding site by exploring a large variety of possible R-groups that would fit in the additional binding pocket present in the new crystal structure. For instance, the algorithm explored an oxazolidinone substituent at the ortho position (Scheme 1, compound **3**) that is able to fit into the new crystal structure but would clash with the Ile223 in the original crystal structure (Figure 7). The FEP+ predicted potency for both enantiomers of the oxazolidinone-substituted

compound showed a distinct difference in potency for both enantiomers; pred. pIC_{50} = 8.96 for (S)-**3**, and pred. pIC_{50} = 5.10 for (R)-**3**. This potency of the racemic compound was confirmed experimentally after synthesis, and evaluation (**3**); pIC_{50} = 6.31




Figure 7. (A) Superimposition of the crystal structure originally used for FEP+ modeling at project initiation (blue) with the crystal structure obtained for compound **1** bound to the receptor (cyan). (B) Docked pose of compound (S)-**3** showing clashes with several residues of the original crystal structure. (C) Docked pose of compound (S)-**3** indicating that the compound fits into the crystal structure of compound **1** bound to the receptor without major clashes.

Different crystal structures of the same protein will be able to accommodate ligands of different sizes and R-group vectors. With the help of AutoDesigner we are able to quickly probe the binding site of every crystal structure obtained for the target protein through the generation of large libraries of project-relevant chemical matter.

Introduction of Polarity via Novel Water Mediated Interactions

Besides fully exploring the protein's binding pocket, and exploring billions of compounds that have the potential to fit the desired physicochemical property space, such large-scale exploration of chemical space also has the added benefit of discovering unexpected and non-intuitive interactions with the receptor. The region of the DAO binding site that interacts with the tail section of the ligand is primarily hydrophobic, but to improve physicochemical properties we were interested to see if more polar groups could be accommodated in this region without sacrificing potency.

In this particular example, one of the final compounds generated by AutoDesigner contains a benzothiadiazole R-group (Figure 8). The compound was predicted to have improved potency (FEP+ $pIC_{50} = 8.05$), and improved LLE (6.35) while at the same time maintaining an excellent CNS MPO score (4.99). Though there are no additional hydrogen bond interactions with the protein, the predicted potency increases as both nitrogen atoms of the benzothiadiazole engage in a water network and the sulfur atom makes additional hydrophobic interactions. The compound was synthesized and assayed, and, at that stage of the project, compound **4** was one of the most potent compounds in the series ($pIC_{50} = 7.41$). The tight integration of FEP+ into the AutoDesigner algorithm enables us to exhaustively explore chemical space and predict the most potent compounds, leading to the discovery of novel chemotypes with unexpected interactions and potency.




Figure 8. FEP+ snapshot showing the binding pose and novel water-mediated interactions of compound **4** with the receptor.

Multiparameter Optimization of Multiple Experimental and Computational Endpoints

Crucial for the lead optimization stage of the project was to design compounds that are not only potent, fit within the desired physicochemical property space and CNS MPO, but at the same time exhibit the desired ADME properties, particularly mouse IV plasma clearance ($\text{Cl}_p < 40$ mL/min/Kg) and brain exposure ($\text{K}_{p,uu} > 0.1$). In this AutoDesigner run, compound **5** was used as the starting ligand, consisting of a pyrazine dione head piece, a thioether linker, and a trifluoromethyl phenyl tail.

Following the final filtering cascade, as described in Figure 2, over 15,000 unique, chemically stable, and drug-like ligands were generated that passed all filters and fit the binding pocket (See Table 2).

In order to enrich for compounds with the requisite properties, machine learning models were generated to predict mouse IV plasma clearance (Cl_p), and the unbound partition coefficient ($K_{p,uu}$) using Schrödinger AutoQSAR and available project data. These machine learning models were applied to the AutoDesigner output and act as additional filtering stages at the end of the algorithm depicted in Figure 2. In total, 12,521 compounds had a predicted $Cl_p < 40$ and a predicted $K_{p,uu} > 0.1$. Additionally, compounds were evaluated on their closest Tanimoto similarity to competitor compounds (<0.5), CNS MPO score (>4), and FEP+ predicted potency. One of the AutoDesigner compounds that was synthesized based on these strict criteria was verified experimentally and met all the desired endpoints was compound **6** (Table 4).

Table 4. Property profile of compound **6**

hDAO pIC_{50}	7.06 (7.77) ^a
Closest Tanimoto similarity to competitor compounds	0.4
LLE	5.7
Pfizer CNS MPO	4.90
Cl_p	26 (<40)
$K_{p,uu}$	0.09 (>0.1)

^a FEP+ predicted pIC_{50} in parentheses.

In order to produce 12,521 ligands that satisfy the desired mouse IV plasma clearance (Cl_p), and unbound partition coefficient ($K_{p,uu}$) criteria, 353 million compounds were evaluated. Taking into account additional project criteria such as low Tanimoto similarity to competitor compounds, FEP+ predicted potency, LLE, and CNS MPO score, the final number of compounds that meet all of the desired criteria is much lower. These results highlight the

challenges of the drug discovery process, and the need for ultra-large-scale exploration of chemical space in order to find the proverbial “needle in the haystack”.

Exploration of Novel DAO Subpocket to Improve Potency

Based on the relatively limited SAR exploration of the DAO subpocket described in the literature, it was postulated that larger groups cannot be accommodated by the relatively narrow subpocket. Ligands with additional substitution on the phenyl ring suffered from a significant decrease in potency.³⁵ We decided to use the AutoDesigner algorithm to quickly and thoroughly explore whether this hypothesis was correct, and what type of groups could potentially fit within the subpocket.

As a starting point for the AutoDesigner algorithm we selected compound **7**, consisting of the pyrazine dione head group linked to a fused bicyclic ring system, and decided to explore the 7-position of the ring system. Introduction of the fused 2,3-dihydro-1,4-benzoxathiine ring led to a gain in potency compared to the uncyclized ligand (**5**) due to stabilization of the linker and tail piece conformation. The AutoDesigner algorithm designed and evaluated 199 million unique ligands, and ultimately 156 ligands passed the final filtering cascade, suggesting that this is a relatively difficult design problem. Even though the subpocket is relatively narrow, there are some compounds that were able to dock into this new subpocket. Not only was the AutoDesigner algorithm able to design compounds that fit all desired criteria, and fit in the new subpocket, but also FEP+ predictions indicated that potency can be improved by exploration of the new subpocket (Figure 9).

Synthesis of one of the most promising AutoDesigner compounds in this series, compound **8**, confirmed the predicted potency and challenged the hypothesis that substitution of the ring system reduces potency. This supports the argument that the AutoDesigner algorithm, and its ability to explore billions of compounds, is beneficial to understand the SAR of the

binding pocket, and is a robust and efficient ligand design tool. Additionally, this produced the most potent compound in the project at the time ($pIC_{50} = 8.11$).




Figure 9. Side by side comparison of surfaces occupied by a publicly disclosed crystallized inhibitor³⁵ (left, PDB: 3W4K), and compound **8** (right) in the DAO binding pocket.

Conclusions

We have recently seen improvements in compute power simultaneously matched by increases in predictive accuracy of computational methods. This has created the opportunity for accurate computational exploration and profiling of chemical space at a scale far larger than possible via traditional experimental approaches. To fully capitalize on these advances, we developed AutoDesigner, which employs a broad search algorithm and cloud computing to carry out successive rounds of chemical space exploration and filtering. We demonstrated the utility of AutoDesigner by applying it towards the design of DAO inhibitors. Over the course of the drug-discovery campaign, over one billion compounds were explored by AutoDesigner using only three ligands as starting points. This highlights the ability to effectively use AutoDesigner even in the absence of large sets of experimental data. Depending on the series being explored

and the specific MPO criteria, it was observed that the size and nature of the chemical space could vary widely. This data helped the team strategically prioritize which compounds to pursue for synthesis.

Synthesis and assay of AutoDesigner compound **2** demonstrate the ability of the algorithm to generate molecules that meet the $<1 \mu\text{M}$ potency cutoff while maintaining the required physicochemical property space. Compound **6** illustrates the ability of AutoDesigner to generate molecules that were not only synthetically tractable and potent, but also met a variety of desired MPO criteria. Compounds **4** and **8** generated by AutoDesigner explore novel interactions and a novel subpocket, respectively. These examples underscore the advantage of being able to rapidly evaluate large amounts of chemical space during lead optimization to help discover molecules that meet multiple project criteria. To date, across a wide variety of drug-discovery projects at Schrödinger, over 500 billion compounds have been explored using AutoDesigner, resulting in the rapid discovery of key compounds that have solved several project issues, and a rapid acceleration of the hit-to-lead and lead optimization processes, as exemplified by the DAO program described in this manuscript.

We have also found that AutoDesigner is most effective when run in close consultation with medicinal chemists. The first reason for this is that although the algorithm is designed to be synthetically aware at each chemical exploration stage (i.e., using previously published transformations, retrosynthetic enumeration, and small R-group decorations), the resulting molecules usually occupy a wide range of synthetic tractability ranging from simple to difficult. This is not a limitation of AutoDesigner per se, as the computational prediction of synthetic feasibility at scale is still an area of active research.³⁶ Nonetheless, in practice this means that collaboration with medicinal chemists to prioritize AutoDesigner compounds for synthesis is essential. Second, medicinal chemists play a key role in defining the permissible property space ranges for the AutoDesigner compounds. As noted above, one of the strengths of AutoDesigner is that the physicochemical ranges are fully customizable, allowing for design beyond the simple

“rule of 5” cutoffs.³⁷ However, delineating these project specific ranges, which often change over the course of the program, benefits from the experience and analytic capabilities of experienced medicinal chemists to ensure that compounds with reasonable properties are designed. Thus, by working collaboratively, computational and medicinal chemists can utilize AutoDesigner to rapidly advance SBDD programs for challenging targets.

Acknowledgments

The authors thank Alexandre Côté, Aleksey Gerasyuto, Jeremy Greenwood, Xianhai Huang, Hideyuki Igawa, Daigo Inoyama, Kyle Konze, Adam Levinson, Andrew Placzek, Hidenori Takahashi, Michael Trzoss, Jiashi Wang, and Jiayi Xu for visual examination of AutoDesigner results and useful feedback on filters for chemical stability, synthetic feasibility and medicinal chemistry liabilities. We would also like to thank René Kanters and Ivan Tubert-Brohman for their contributions to the Python codebase of AutoDesigner.

Abbreviations used

AL-FEP Active-Learning FEP; CNS MPO central nervous system multiparameter optimization; DAO D-amino acid oxidase; FAD flavin adenine dinucleotide; FEP free energy perturbation; H2L hit-to-lead; hDAO human DAO; HPC high performance computing; LO lead optimization; MPO multiparameter optimization; NMDA N-methyl D-aspartate; PK pharmacokinetics; QSAR quantitative structure-activity relationship; SAR structure-activity relationship; SFI Solubility Forecast Index; SMARTS SMILES arbitrary target specification; SMILES simplified molecular-input line-entry system; TPP target product profile; TPSA topological polar surface area.

References

- (1) Schneider, G. Automating Drug Discovery. *Nat. Rev. Drug Discov.* **2018**, *17* (2), 97–113.
- (2) Hann, M. M. Molecular Obesity, Potency and Other Addictions in Drug Discovery. *Medchemcomm* **2011**, *2* (5), 349–355.
- (3) Schneider, G. Mind and Machine in Drug Design. *Nat Mach Intell* **2019**, *1* (3), 128–130.
- (4) Zhou, Z.; Kearnes, S.; Li, L.; Zare, R. N.; Riley, P. Optimization of Molecules via Deep Reinforcement Learning. *Sci. Rep.* **2019**, *9* (1), 10752.
- (5) Green, D. V. S.; Pickett, S.; Luscombe, C.; Senger, S.; Marcus, D.; Meslamani, J.; Brett, D.; Powell, A.; Masson, J. BRADSHAW: A System for Automated Molecular Design. *J. Comput. Aided Mol. Des.* **2020**, *34* (7), 747–765.
- (6) Besnard, J.; Ruda, G. F.; Setola, V.; Abecassis, K.; Rodriguiz, R. M.; Huang, X.-P.; Norval, S.; Sassano, M. F.; Shin, A. I.; Webster, L. A.; Simeons, F. R. C.; Stojanovski, L.; Prat, A.; Seidah, N. G.; Constam, D. B.; Bickerton, G. R.; Read, K. D.; Wetsel, W. C.; Gilbert, I. H.; Roth, B. L.; Hopkins, A. L. Automated Design of Ligands to Polypharmacological Profiles. *Nature* **2012**, *492* (7428), 215–220.
- (7) Wang, L.; Wu, Y.; Deng, Y.; Kim, B.; Pierce, L.; Krilov, G.; Lypyan, D.; Robinson, S.; Dahlgren, M. K.; Greenwood, J.; Romero, D. L.; Masse, C.; Knight, J. L.; Steinbrecher, T.; Beuming, T.; Damm, W.; Harder, E.; Sherman, W.; Brewer, M.; Wester, R.; Murcko, M.; Frye, L.; Farid, R.; Lin, T.; Mobley, D. L.; Jorgensen, W. L.; Berne, B. J.; Friesner, R. A.; Abel, R. Accurate and Reliable Prediction of Relative Ligand Binding Potency in Prospective Drug Discovery by Way of a Modern Free-Energy Calculation Protocol and Force Field. *J. Am. Chem. Soc.* **2015**, *137* (7), 2695–2703.
- (8) Shim, S. S.; Hammonds, M. D.; Kee, B. S. Potentiation of the NMDA Receptor in the Treatment of Schizophrenia: Focused on the Glycine Site. *Eur. Arch. Psychiatry Clin. Neurosci.* **2008**, *258* (1), 16–27.
- (9) Yang, C. R.; Svensson, K. A. Allosteric Modulation of NMDA Receptor via Elevation of Brain Glycine and D-Serine: The Therapeutic Potentials for Schizophrenia. *Pharmacol. Ther.* **2008**, *120* (3), 317–332.
- (10) Fujita, Y.; Ishima, T.; Hashimoto, K. Supplementation with D-Serine Prevents the Onset of Cognitive Deficits in Adult Offspring after Maternal Immune Activation. *Sci. Rep.* **2016**, *6*, 37261.
- (11) Hagiwara, H.; Iyo, M.; Hashimoto, K. Neonatal Disruption of Serine Racemase Causes Schizophrenia-like Behavioral Abnormalities in Adulthood: Clinical Rescue by D-Serine. *PLoS One* **2013**, *8* (4), e62438.
- (12) Sacchi, S.; Rosini, E.; Pollegioni, L.; Molla, G. D-Amino Acid Oxidase Inhibitors as a Novel Class of Drugs for Schizophrenia Therapy. *Curr. Pharm. Des.* **2013**, *19* (14), 2499–2511.
- (13) Ferraris, D. V.; Tsukamoto, T. Recent Advances in the Discovery of D-Amino Acid Oxidase Inhibitors and Their Therapeutic Utility in Schizophrenia. *Current Pharmaceutical Design*. 2011, pp 103–111. <https://doi.org/10.2174/138161211795049633>.
- (14) Tang, H.; Jensen, K.; Houang, E.; McRobb, F.; Bhat, S.; Svensson, M.; Bochevarov, A.; Dahlgren, M.; Bell, J.; Day, T.; Frye, L.; Skene, R.; Lewis, J.; Osborne, J.; Tierney, J.; Gordon, J.; Palomero-Vazquez, M.; Gallati, C.; Chapman, R.; Jones, D.; Hirst, K.; Sephton, M.; Chauhan, A.; Sharpe, A.; Tardia, P.; Dechaux, E.; Taylor, A.; Waddell, R.; Valentine, A.; Janssens, H.; Aziz, O.; Bloomfield, D.; Ladha, S.; Fraser, I.; Ellard, J. Discovery of a Novel Class of D-Amino Acid Oxidase (DAO) Inhibitors with the Schrödinger Computational Platform, 2021. <https://doi.org/10.33774/chemrxiv-2021-dkf1k>.
- (15) Konze, K. D.; Bos, P. H.; Dahlgren, M. K.; Leswing, K.; Tubert-Brohman, I.; Bortolato, A.; Robbason, B.; Abel, R.; Bhat, S. Reaction-Based Enumeration, Active Learning, and Free Energy Calculations To Rapidly Explore Synthetically Tractable Chemical Space and Optimize Potency of Cyclin-Dependent Kinase 2 Inhibitors. *J. Chem. Inf. Model.* **2019**, *59* (9), 3782–3793.
- (16) Ghanakota, P.; Bos, P. H.; Konze, K. D.; Staker, J.; Marques, G.; Marshall, K.; Leswing, K.; Abel, R.; Bhat, S. Combining Cloud-Based Free-Energy Calculations, Synthetically Aware Enumerations, and Goal-Directed Generative Machine Learning for Rapid Large-Scale Chemical Exploration and Optimization. *J. Chem. Inf. Model.* **2020**, *60* (9), 4311–4325.
- (17) Knight, J. L.; Leswing, K.; Bos, P. H.; Wang, L. Impacting Drug Discovery Projects with Large-Scale Enumerations, Machine Learning Strategies, and Free-Energy Predictions. In *Free Energy Methods in Drug Discovery: Current State and Future Directions*; ACS Symposium Series;

- American Chemical Society, 2021; Vol. 1397, pp 205–226.
- (18) Kim, S.; Chen, J.; Cheng, T.; Gindulyte, A.; He, J.; He, S.; Li, Q.; Shoemaker, B. A.; Thiessen, P. A.; Yu, B.; Zaslavsky, L.; Zhang, J.; Bolton, E. E. PubChem in 2021: New Data Content and Improved Web Interfaces. *Nucleic Acids Res.* **2021**, *49* (D1), D1388–D1395.
 - (19) Gaulton, A.; Hersey, A.; Nowotka, M.; Bento, A. P.; Chambers, J.; Mendez, D.; Mutowo, P.; Atkinson, F.; Bellis, L. J.; Cibrián-Uhalte, E.; Davies, M.; Dedman, N.; Karlsson, A.; Magariños, M. P.; Overington, J. P.; Papadatos, G.; Smit, I.; Leach, A. R. The ChEMBL Database in 2017. *Nucleic Acids Res.* **2017**, *45* (D1), D945–D954.
 - (20) Dalke, A.; Hert, J.; Kramer, C. Mmpdb: An Open-Source Matched Molecular Pair Platform for Large Multiproperty Data Sets. *J. Chem. Inf. Model.* **2018**, *58* (5), 902–910.
 - (21) Mey, A. S. J. S.; Allen, B. K.; Macdonald, H. E. B.; Chodera, J. D.; Hahn, D. F.; Kuhn, M.; Michel, J.; Mobley, D. L.; Naden, L. N.; Prasad, S.; Rizzi, A.; Scheen, J.; Shirts, M. R.; Tresadern, G.; Xu, H. Best Practices for Alchemical Free Energy Calculations [Article v1.0]. *Living J Comput Mol Sci* **2020**, *2* (1). <https://doi.org/10.33011/livecoms.2.1.18378>.
 - (22) Cappel, D.; Jerome, S.; Hessler, G.; Matter, H. Impact of Different Automated Binding Pose Generation Approaches on Relative Binding Free Energy Simulations. *J. Chem. Inf. Model.* **2020**, *60* (3), 1432–1444.
 - (23) Dixon, S. L.; Smodyrev, A. M.; Rao, S. N. PHASE: A Novel Approach to Pharmacophore Modeling and 3D Database Searching. *Chem. Biol. Drug Des.* **2006**, *67* (5), 370–372.
 - (24) Dixon, S. L.; Smodyrev, A. M.; Knoll, E. H.; Rao, S. N.; Shaw, D. E.; Friesner, R. A. PHASE: A New Engine for Pharmacophore Perception, 3D QSAR Model Development, and 3D Database Screening: 1. Methodology and Preliminary Results. *J. Comput. Aided Mol. Des.* **2006**, *20* (10-11), 647–671.
 - (25) Sastry, G. M.; Adzhigirey, M.; Day, T.; Annabhimoju, R.; Sherman, W. Protein and Ligand Preparation: Parameters, Protocols, and Influence on Virtual Screening Enrichments. *J. Comput. Aided Mol. Des.* **2013**, *27* (3), 221–234.
 - (26) Shelley, J. C.; Cholleti, A.; Frye, L. L.; Greenwood, J. R.; Timlin, M. R.; Uchimaya, M. Epik: A Software Program for pK(a) Prediction and Protonation State Generation for Drug-like Molecules. *J. Comput. Aided Mol. Des.* **2007**, *21* (12), 681–691.
 - (27) Bochevarov, A. D.; Harder, E.; Hughes, T. F.; Greenwood, J. R.; Braden, D. A.; Philipp, D. M.; Rinaldo, D.; Halls, M. D.; Zhang, J.; Friesner, R. A. Jaguar: A High-performance Quantum Chemistry Software Program with Strengths in Life and Materials Sciences. *International Journal of Quantum Chemistry*. 2013, pp 2110–2142. <https://doi.org/10.1002/qua.24481>.
 - (28) Chen, W.; Deng, Y.; Russell, E.; Wu, Y.; Abel, R.; Wang, L. Accurate Calculation of Relative Binding Free Energies between Ligands with Different Net Charges. *J. Chem. Theory Comput.* **2018**, *14* (12), 6346–6358.
 - (29) Roos, K.; Wu, C.; Damm, W.; Reboul, M.; Stevenson, J. M.; Lu, C.; Dahlgren, M. K.; Mondal, S.; Chen, W.; Wang, L.; Abel, R.; Friesner, R. A.; Harder, E. D. OPLS3e: Extending Force Field Coverage for Drug-Like Small Molecules. *J. Chem. Theory Comput.* **2019**, *15* (3), 1863–1874.
 - (30) Ross, G. A.; Russell, E.; Deng, Y.; Lu, C.; Harder, E. D.; Abel, R.; Wang, L. Enhancing Water Sampling in Free Energy Calculations with Grand Canonical Monte Carlo. *J. Chem. Theory Comput.* **2020**, *16* (10), 6061–6076.
 - (31) Dixon, S. L.; Duan, J.; Smith, E.; Von Bargen, C. D.; Sherman, W.; Repasky, M. P. AutoQSAR: An Automated Machine Learning Tool for Best-Practice Quantitative Structure-Activity Relationship Modeling. *Future Med. Chem.* **2016**, *8* (15), 1825–1839.
 - (32) Central Nervous System Multiparameter Optimization Desirability: Application in Drug Discovery. <https://doi.org/10.1021/acscchemneuro.6b00029.s001>.
 - (33) Pajouhesh, H.; Lenz, G. R. Medicinal Chemical Properties of Successful Central Nervous System Drugs. *NeuroRx* **2005**, *2* (4), 541–553.
 - (34) Leung, C. S.; Leung, S. S. F.; Tirado-Rives, J.; Jorgensen, W. L. Methyl Effects on Protein-Ligand Binding. *J. Med. Chem.* **2012**, *55* (9), 4489–4500.
 - (35) Hondo, T.; Warizaya, M.; Niimi, T.; Namatame, I.; Yamaguchi, T.; Nakanishi, K.; Hamajima, T.; Harada, K.; Sakashita, H.; Matsumoto, Y.; Orita, M.; Takeuchi, M. 4-Hydroxypyridazin-3(2H)-One Derivatives as Novel D-Amino Acid Oxidase Inhibitors. *J. Med. Chem.* **2013**, *56* (9), 3582–3592.
 - (36) Thakkar, A.; Chadimová, V.; Bjerrum, E. J.; Engkvist, O.; Reymond, J.-L. Retrosynthetic Accessibility Score (RAscore) - Rapid Machine Learned Synthesizability Classification from AI Driven

Retrosynthetic Planning. *Chem. Sci.* **2021**, *12* (9), 3339–3349.

- (37) Lipinski, C. A.; Lombardo, F.; Dominy, B. W.; Feeney, P. J. Experimental and Computational Approaches to Estimate Solubility and Permeability in Drug Discovery and Development Settings. *Adv. Drug Deliv. Rev.* **2001**, *46* (1-3), 3–26.

Supplementary Material

S1

Comparison of transforms (SMIRKS) extracted from ChEMBL and PubChem

Total number of unique reaction SMARTS:

PubChem: 290,869,841 (291 M).

ChEMBL: 25,725,737 (26 M).

Overlap/difference:

Intersection: reaction SMARTS common to both: 6,283,961 (6.3 M).

Difference: reaction SMARTS in PubChem but not in ChEMBL: 284,585,880 (285 M).

PubChem:

Number of redundant transforms: 91,669,639,245 (92 Billion).

Number of unique transforms: 290,869,841 (291 Million).

ChEMBL:

Number of redundant transforms: 7,161,388,428 (7.2 Billion).


Number of unique transforms: 25,725,737 (26 Million).

Table S1. Classification criteria. These properties are discussed below. Note that NCC, NR, and HAC correspond to the number of chiral centers, number of rings, and number of heavy atoms respectively. Structures are assigned to a class by the last successful match on all criteria, proceeding from right to left.

Near drug-like	Drug-like	Lead-like	Fragment
$-1.5 \leq \text{AlogP} \leq 5.5$	$-1 \leq \text{AlogP} \leq 4$	$0 \leq \text{AlogP} \leq 3$	$\text{AlogP} \leq 3$
$150 \leq \text{MW} \leq 575$	$250 \leq \text{MW} \leq 500$	$250 \leq \text{MW} \leq 375$	$\text{MW} > 110$
$30 < \text{PSA} < 150$	$50 < \text{PSA} < 130$	$\text{PSA} < 110$	$\text{PSA} \leq 110$
$\text{HBD} \leq 5$	$\text{HBD} \leq 5$	$\text{HBD} \leq 2$	$\text{HBD} \leq 3$
$\text{HBA} \leq 12$	$\text{HBA} \leq 10$	$\text{HBA} \leq 5$	$\text{HBA} \leq 5$
$\text{RB} \leq 10$	$\text{RB} \leq 10$	$\text{RB} \leq 10$	$\text{RB} \leq 3$
$\text{NCC} \leq 3$	$\text{NCC} \leq 3$	$\text{NCC} \leq 1$	
			$\text{NR} \geq 1$
			$\text{HAC} \leq 18$

S2 Details on the AutoDesigner starting from compound 1

Table S2a. Number of compounds generated in the AutoDesigner run from compound 1.^{a,b}



Input Ligand	MMP	Trim 1	PF	Trim 2	Deco 1	Deco 2	Total ^c	Final Output
1	157M	24M	466M	24M	44M	45M	761M	515K

^a Immutable region/common core of the starting ligands highlighted in blue. ^b Abbreviations used for generative stages: MMP: matched molecular pair transformations; Trim: recursive trimming; PF: PathFinder reaction-based enumeration; Deco: R-group decoration. ^c Total number of unique compounds explored.

Table S2b. Compound property space of the AutoDesigner output from compound 1.

Property	Value
Ligand total charge	0
Ligand absolute total charge	0
Molecular Weight	0 - 350
Fraction sp ³	0 - 1
Topological Polar Surface Area	70 - 110
AlogP	0 - 3
Number of Rotatable Bonds	0 - 8
Maximum number of Consecutive Rotatable Bonds	0 - 4
Chiral Centers	0 - 1
Number of Rings	0 - 10

Number of Aromatic Rings	0 - 3
Number of Saturated Rings	0 - 10
Number of Spiro atoms	0 - 1
Number of Bridgehead atoms	0 - 2
Number of Hydrogen Bond Donors	0 - 5
Number of Hydrogen Bond Acceptors	0 - 10
Maximum Ring Size	7
Maximum number of Sulfur Atoms	1
Maximum number of Alkynes	0
Number of Quaternary Substituted Carbon Atoms	1

Table S2c. Classification of the final output starting from compound **1**

	Drug-like	Lead-like	Near drug-like	Fragment	Others
Count	385,198	129,151	85	221	3
Percentage	74.8%	25.1%	0.0%	0.0%	0.0%

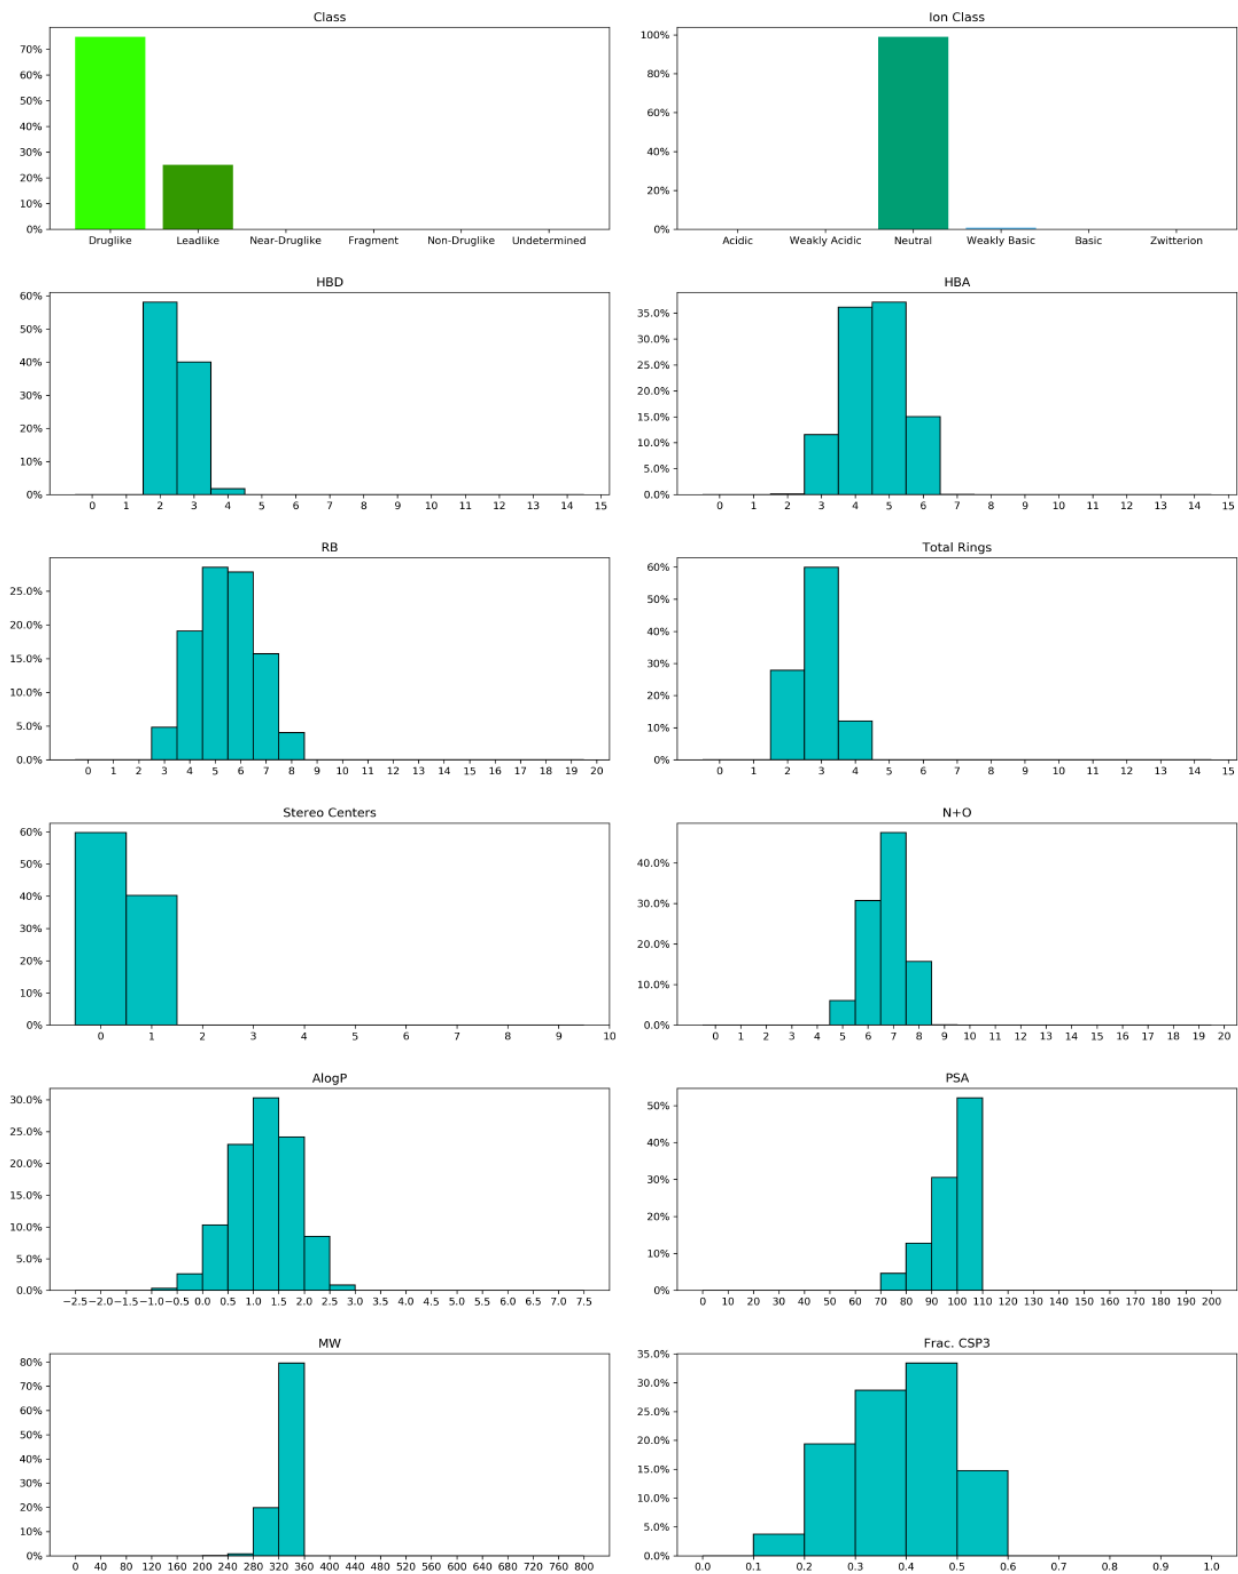



Figure S2d. Property distributions of selected properties from the AutoDesigner run starting from compound 1.

S3 Details on the AutoDesigner starting from compound 5

Table S3a. Number of compounds generated in the AutoDesigner run from compound 5.^{a,b}



Input Ligand	MMP	Trim 1	PF	Trim 2	Deco 1	Deco 2	Total ^c	Final Output
5	101M	121M	234M	2M	2M	2M	353M	15K

^a Immutable region/common core of the starting ligands highlighted in blue. ^b Abbreviations used for generative stages: MMP: matched molecular pair transformations; Trim: recursive trimming; PF: PathFinder reaction-based enumeration; Deco: R-group decoration. ^c Total number of unique compounds explored.

Table S3b. Compound property space of the AutoDesigner output from compound 5.

Property	Value
Ligand total charge	0
Ligand absolute total charge	0
Molecular Weight	0 - 325
Fraction sp ³	0 - 1
Topological Polar Surface Area	70 - 110
AlogP	0 - 3.5
Number of Rotatable Bonds	0 - 8
Maximum number of Consecutive Rotatable Bonds	0 - 3
Chiral Centers	0 - 1
Number of Rings	0 - 6

Number of Aromatic Rings	0 - 3
Number of Saturated Rings	0 - 3
Number of Spiro atoms	0 - 1
Number of Bridgehead atoms	0 - 2
Number of Hydrogen Bond Donors	0 - 2
Number of Hydrogen Bond Acceptors	0 - 8
Maximum Ring Size	7
Maximum number of Sulfur Atoms	1
Maximum number of Alkynes	0
Number of Quaternary Substituted Carbon Atoms	1

Table S3c. Classification of the final output starting from compound **5**.

	Drug-like	Lead-like	Near drug-like	Fragment	Others
Count	2,040	12,474	0	162	0
Percentage	13.9%	85.0%	0.0%	1.1%	0.0%

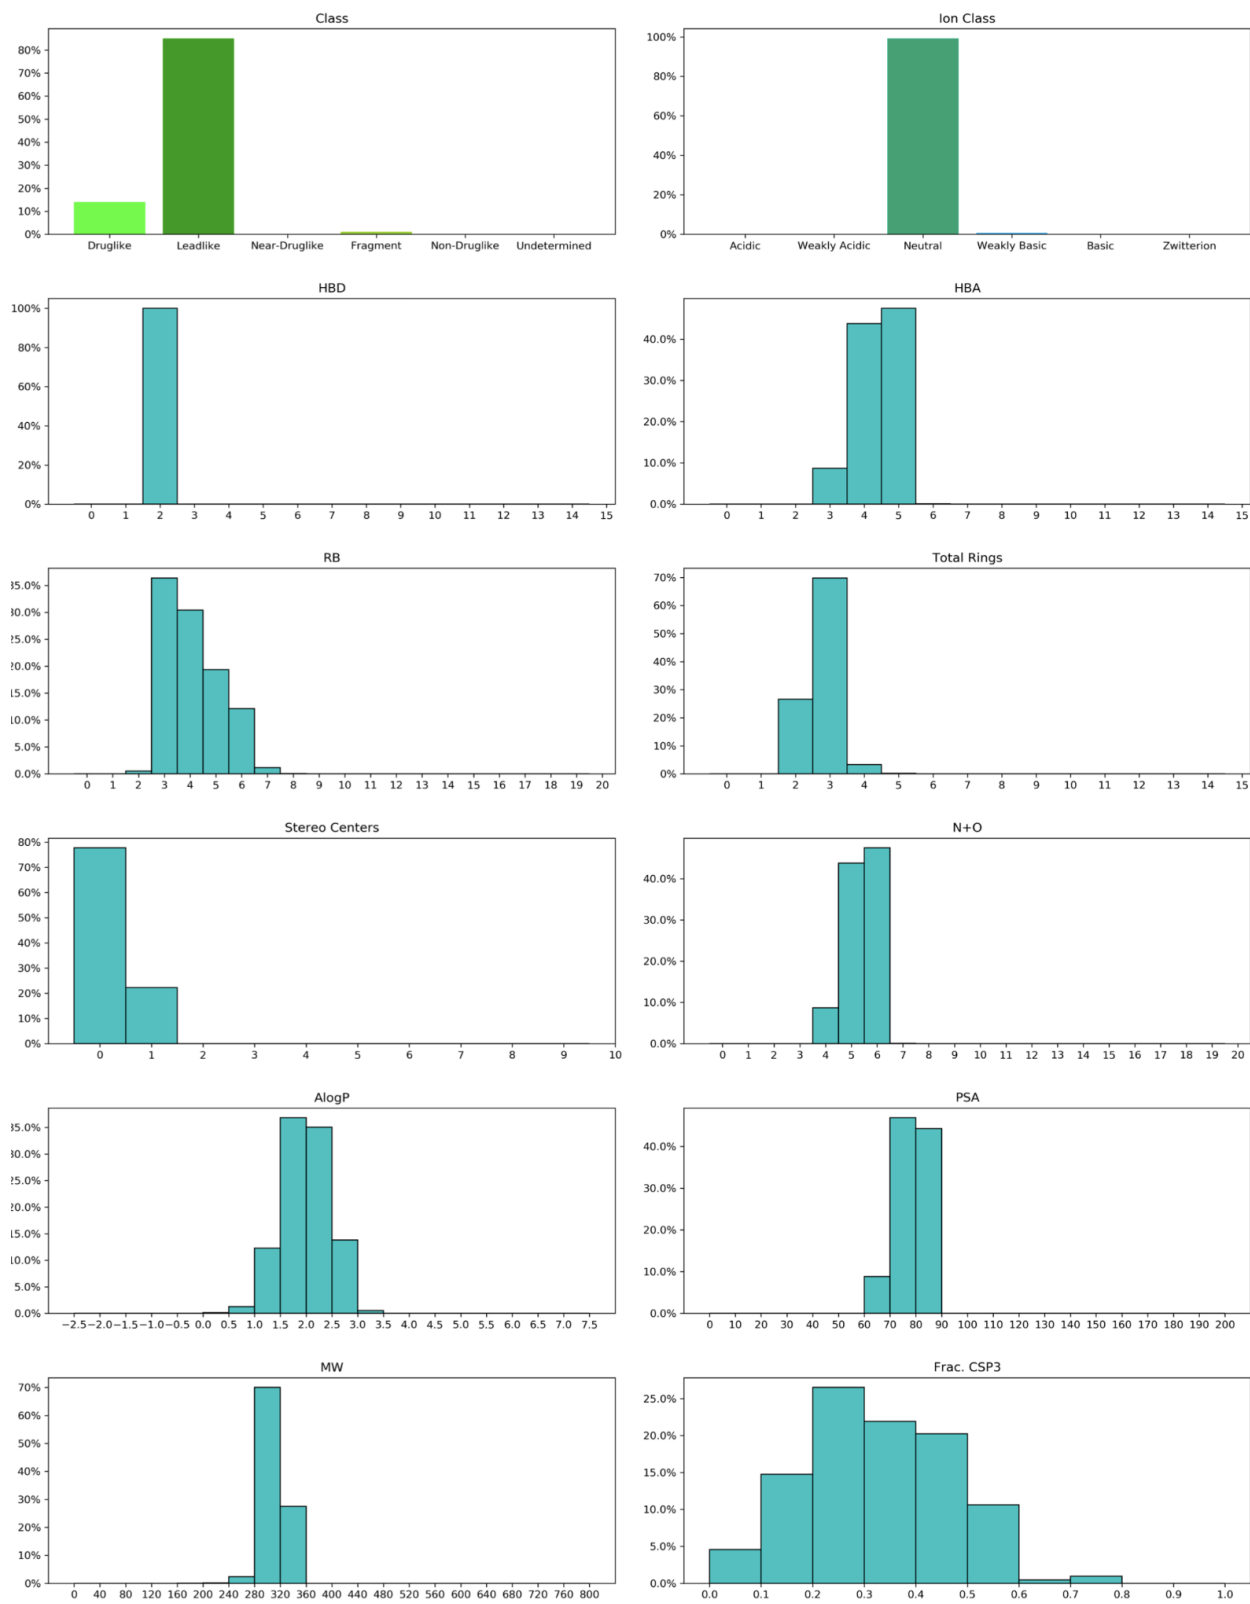



Figure S3d. Property distributions of selected properties from the AutoDesigner run starting from compound 5.

S4 Details on the AutoDesigner starting from compound 7

Table S4a. Number of compounds generated in the AutoDesigner run from compound 7.^{a,b}



Input Ligand	MMP	Trim 1	PF	Trim 2	Deco 1	Deco 2	Total ^c	Final Output
7	161M	13M	23M	458K	348K	62K	199M	156

^a Immutable region/common core of the starting ligands highlighted in blue. ^b Abbreviations used for generative stages: MMP: matched molecular pair transformations; Trim: recursive trimming; PF: PathFinder reaction-based enumeration; Deco: R-group decoration. ^c Total number of unique compounds explored.

Table S4b. Compound property space of the AutoDesigner output from compound 7.

Property	Value
Ligand total charge	0
Ligand absolute total charge	0
Molecular Weight	0 - 401
Fraction sp ³	0 - 1
Topological Polar Surface Area	70 - 110
AlogP	0 - 4
Number of Rotatable Bonds	0 - 8
Maximum number of Consecutive Rotatable Bonds	0 - 3
Chiral Centers	0 - 1
Number of Rings	0 - 10

Number of Aromatic Rings	0 - 3
Number of Saturated Rings	0 - 10
Number of Spiro atoms	0 - 1
Number of Bridgehead atoms	0 - 2
Number of Hydrogen Bond Donors	0 - 2
Number of Hydrogen Bond Acceptors	0 - 10
Maximum Ring Size	7
Maximum number of Sulfur Atoms	1
Maximum number of Alkynes	0
Number of Quaternary Substituted Carbon Atoms	1

Table S4c. Classification of the final output starting from compound 7.

	Drug-like	Lead-like	Near drug-like	Fragment	Others
Count	80	75	0	1	0
Percentage	51.3%	48.1%	0.0%	0.6%	0.0%




Figure S4d. Property distributions of selected properties from the AutoDesigner run starting from compound 7.